Birational maps

Merrick Cai

March 10, 2024

Contents

1 Overview

1.1 Introduction

The goal of this note is to summarize [Har13, V.5], birational transformations. I also reference [Bea96], specifically with regards to the proof of [Bea96, Theorem II.7], "elimination of indeterminancy." The main goal of this note is to explain [Har13, V.5], so I will mostly follow that subchapter.

1.2 Main results

The main results are detailed in §3. For reference, the two main results are:

Theorem 1.1 (Factorization of birational maps of surfaces, 3.1): Let $T : X' \dashrightarrow X$ be a birational map of surfaces. Then T can be written as the composition of blowups (at points) and their inverses. Explicitly, there exists a surface S and maps $\eta : S \to X'$, $\psi : S \to X$ such that both η, ψ are a finite sequence of blowups, and they make the following diagram commute:

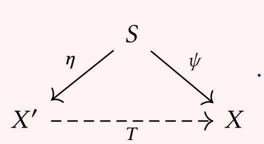

This allows us to describe every birational map of surfaces via blowups.

Theorem 1.2 (Castelnuovo's criterion, 3.3): Let Y be a curve on a (smooth projective) surface X, such that $Y \cong \mathbb{P}^1$ and $Y^2 = -1$. Then $X \cong \mathrm{Bl}_p X_0$ for some (smooth projective) surface X_0, such that Y is the exceptional curve of this blowup.

This allows us to understand exactly when we can write surfaces as blowups of other surfaces.

1.3 Summary

Let me summarize the basic ideas that are going on. We wish to study birational maps, which are not defined everywhere, hence are rather difficult to work with. Theorem 2.14 allows us to regard rational maps from surfaces as two genuine morphisms, albeit with one inverted. This is very useful because we now can replace a birational map, which is only defined on an open set of the source, with maps which are genuine morphisms of projective varieties. From here, will see that every birational morphism of surfaces is actually a sequence of blowups, which further adds to our explicit understanding of birational maps of surfaces - now we understand them as a sequence of blowups and their inverses. So the theory of birational maps of surfaces boils down to blowing up points. We can always blow up; it only remains to ask when we can blow *down*, i.e., the surface is the blowup of another surface. This will be answered precisely by Castelnuovo's theorem (3.3), which will complete the general picture.

2 Setup

2.1 Definitions

Let X and Y be projective varieties.

Definition 2.1: A **rational map** $T : X \dashrightarrow Y$ is the data of a dense open subset $U \subset X$ and a morphism $\varphi : U \to Y$.

Definition 2.2: A **birational map** $T : X \dashrightarrow Y$ is a rational map such that the morphism $\varphi : U \to Y$ induces an isomorphism on some dense open subset $V \subset U$.

In other words, a birational map is a map which need not be defined on the entire source, only *almost* everywhere, and gives an isomorphism between dense open subsets of the source and the target. Note that the U in this definition is *not* necessarily the open set mapped isomorphically onto its image; it is merely an open set for which T is defined on.

Remark 2.3: An equivalent condition to inducing an isomorphism on a dense open subset is to say that φ_T induces an isomorphism of function fields $K(X) \xrightarrow{\sim} K(Y)$, i.e., an isomorphism on the stalks of the generic points.

Definition 2.4: A **birational morphism** $T : X \to Y$ is a birational map which is a legitimate morphism, i.e., defined everywhere (so we can take U to be all of X).

This may be slightly confusing; a birational *map* may not be defined everywhere, while a birational *morphism* is defined everywhere. The difference in notation will be $\dashrightarrow$ vs. $\to$. Both have to be *birational*, i.e., "generically an isomorphism." A birational map to a normal variety also satisfies the following:

Proposition 2.5: Let $f : X \to Y$ be a birational morphism of projective varieties, with Y normal. Then $f_* \mathcal{O}_X = \mathcal{O}_Y$.

Proof. The problem is local on Y, so assume $Y = \operatorname{Spec} A$ and $X = \operatorname{Spec} B$. Then $f_* \mathcal{O}_X = B$ as an A-algebra (under the map $f^\sharp : A \to B$). Therefore B is a finitely generated A-algebra, and both A, B are integral domains with the same field of fractions (as f is birational). Since Y is normal, A is integrally closed, so $B = A$, and hence $f_* \mathcal{O}_X = \mathcal{O}_Y$. □

Now, a general birational map may not be defined everywhere. If we have (U, φ) and (V, ψ) both representing the birational map $T : X \dashrightarrow Y$, then we can glue them to obtain a map $\gamma : U \cup V \to Y$ representing T. As such, there is a *largest* open set $U \subset X$ for which T is representated by a morphism $U \to Y$.

Definition 2.6: We say that $T : X \dashrightarrow Y$ is **defined at the points of** U. We also say that $X - U$, the points where T cannot be defined, are the **fundamental points** of T.

In particular, birational morphisms have no fundamental points.

Example 2.7: Let $p \in X$, and let $\pi : \mathrm{Bl}_p X \to X$ be the blowup of X at p. Then π is a birational morphism, hence has no fundamental points. On the other hand, $\pi^{-1} : X \dashrightarrow \mathrm{Bl}_p X$ is a birational map, and can be defined everywhere except at $p \in X$. It follows that p is the only fundamental point of the birational map π^{-1}.

Although a birational map $T : X \dashrightarrow Y$ may not be defined at some points in X, we can still assign it a subset of Y. In the case of blowing up a point $\pi : \mathrm{Bl}_p X \to X$, we wish for this subset associated to $p \in X$ under the birational map π^{-1} to be the exceptional divisor E.

Definition 2.8: Let $T : X \dashrightarrow Y$ be a birational map, represented by $\varphi : U \to Y$, where U is the maximal open set for which T is defined. Let $\Gamma_0 \subset U \times Y \subset X \times Y$ be the graph of φ. Then we call the closure $\Gamma := \overline{\Gamma_0} \subset X \times Y$ the **graph of** T; it comes equipped with projection maps $\mathrm{pr}_X, \mathrm{pr}_Y$ to X and Y. To a subset $Z \subset X$, we define the **total transform of** Z to be

$$\widetilde{T}(Z) := \mathrm{pr}_Y(\mathrm{pr}_X^{-1}(Z)).$$

It's easy to check that for points $p \in X$ at which T is defined, then $\widetilde{T}(p) = \varphi(p)$. In other words, the total transform agrees with the map wherever it's already defined

Example 2.9: Let $\pi : \mathrm{Bl}_p X \to X$ be the blowup of X at p. Then π^{-1} has a single fundamental point p. At this point, $\widetilde{\pi^{-1}}(p) = \pi^{-1}(p) = E$,

2.2 Preliminary results

Let's first establish some preliminary results to understand fundamental points and total transforms.

Proposition 2.10: Suppose $T : X \dashrightarrow Y$ is a birational map of projective varieties, and X is normal. Then the fundamental points of T form a closed subset of codimension ≥ 2.

Proof. Since T induces an isomorphism of function fields, it must be defined at the generic point of X. Now any codimension 1 point $p \in X$ necessarily will have its local ring $\mathcal{O}_{X,p}$ be a discrete valuation ring in $K(X)$. By the valuative criterion of properness (4.1), it follows that T must also be defined at P (here we use that Y is projective). □

Corollary 2.11: Let $T : X \dashrightarrow Y$ be a birational map of projective surfaces, with X normal. Then the fundamental points of X are a finite set of points.

We now know that the fundamental points form a "small" subset of the source. We'd also like to understand what the total transform looks like. Wherever T is defined, we know that the total transform is just the image of that subset as a morphism; the question is what happens to the fundamental points.

Proposition 2.12: Let $T : X \to Y$ be a birational transformation of projective varieties, with X normal. If p is a fundamental point of T, then the total transform $\widetilde{T}(p)$ is connected and dimension ≥ 1.

Proof. Let $p \in X$ be a fundamental point and let Γ be the graph of T. Then $\mathrm{pr}_X^{-1}(p) \subset \Gamma$ is mapped isomorphically onto $\widetilde{T}(p)$ by pr_Y. Now $\mathrm{pr}_X : \Gamma \to X$ is a birational projective morphism, so Zariski main theorem (4.3) implies

that $\mathrm{pr}_X^{-1}(p)$ is connected; it follows that so too is $\widetilde{T}(p)$.
It suffices now to show that $\mathrm{pr}_X^{-1}(p)$ is dimension ≥ 1. Suppose not, that it was dimension 0. The key is that upper-semicontinuity of fiber dimensions implies that we must have an open neighborhood $V \ni p$ in X where the fibers under $\mathrm{pr}_X : \Gamma \to X$ are all dimension 0. But V is normal, and $\mathrm{pr}_X : \mathrm{pr}_X^{-1}(V) \to V$ is a finite morphism, so this must be an isomorphism; since T is defined on an open dense subset of X, which must intersect V on an open dense subset, it follows that T must be defined on all of V, contradicting that $p \in V$ is a fundamental point. Therefore by contradiction, we find that $\widetilde{T}(p)$ has dimension ≥ 1. □

Intuitively, this is just saying that if the fiber of a fundamental point was indeed dimension 0, then it must be connected, hence a single point. But the normality of X implies that this single point must also be reduced, so in fact we could have just defined the birational map at this point.

Remark 2.13 (Characterization of fundamental points): Another way to understand fundamental points are *exactly the points for which the total transform is large.* The total transform of a point where T is defined is just a point; the total transform of a fundamental point is a connected, dimension ≥ 1 subvariety.

2.3 Elimination of indeterminancy

This is the main fact which allows us to "factor" a rational map into a morphism and an inverse of another morphism.

Theorem 2.14 (Elimination of indeterminancy): Let $\phi : S \dashrightarrow X$ be a rational map from a surface S to a projective variety X. Then there exists a surface S' such that the diagram

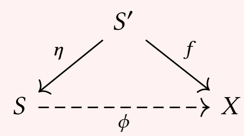

commutes. Furthermore, η is constructed explicitly as a finite sequence of blowups (of points).

Proof. Since X is projective, we have a closed embedding $X \hookrightarrow \mathbb{P}^N$ for some projective space $\mathbb{P}^N$ (and also require that $\phi(S)$ does not lie in a hyperplane of $\mathbb{P}^N$, otherwise just take that hyperplane, which is $\mathbb{P}^{N-1}$). Since ϕ is rational, it's defined on a dense open subset of S. We will see that η is a finite sequence of blowups; in particular, S' has a dense open subset mapped isomorphically under η to a dense open subset of S, and intersecting this with the open set on which ϕ is defined, we have an open subset U in S' and S identified under η, all of which is mapped to $X \subset \mathbb{P}^N$. Since f is continuous, then $f(S')$ must lie in the closure of this open set, which lies in $X \subset \mathbb{P}^N$, so it suffices to just define the map $f : S' \to \mathbb{P}^N$.
Now let's define S'. First, $\phi : U \to X \hookrightarrow \mathbb{P}^N$ is defined by some linear system $P \subset |D|$ with no fixed component (i.e., the base points form a subset of dimension 0) This can be seen again by the valuative criterion of properness (4.1), as ϕ is defined at the generic point, hence at all points of codimension 1; therefore the base points must be codimension 2 in a surface, which is just a finite set of points. If P has no base point, then ϕ is a legitimate morphism, and we can just take $S' = S$. Otherwise, assume that P has a base point $x \in S$.
Then we can take the blowup $\epsilon : S_1 = \mathrm{Bl}_x S \to S$. Now $\epsilon^* P \subset |\epsilon^* D|$ cannot have any fixed components away from $\epsilon^{-1}(x)$, as ϵ is an isomorphism away from $\epsilon^{-1}(x)$, and by hypothesis P had no fixed components. If $\epsilon^* P$ indeed has a fixed component, then it must be the exceptional curve $E = \epsilon^{-1}(x)$, with some multiplicity $m \geq 1$. Then just take the linear system $P_1 := \epsilon^* P - mE \subset |\epsilon^* D - mE|$; this now has no fixed component, and away from the exceptional curve, this induces the exact same map that P did. In other words, P induced a perfectly reasonable map on most of S; where there is an issue (namely, a base point), we simply blow up that point. Now by blowing up, we have the same map induced by $\epsilon^* P$ away from the exceptional curve, so any new issues stay confined to the exceptional locus, which we rectify by just subtracting the fixed-component-multiplicities off without changing the map away

from the exceptional curve.

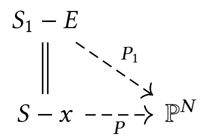

Now we get a rational map $\phi_1 : S_1 \dashrightarrow \mathbb{P}^N$. If this has no base points, we are done. If it does not, we blow up at a base point again. We just need to see that this terminates. The reason that this terminates is because each P_i is a linear system in $|D_n| = |\epsilon_n^* D_{n-1} - m_n E_n|$. But since each P_n has fixed component, then $D_n^2 \geq 0$, but $D_n^2 = D_{n-1}^2 - m_n^2 < D_{n-1}^2$, so this is a strictly decreasing sequence, hence it must terminate.
We therefore eventually arrive at a surface S', which is produced from S by a finite sequence of blowups (at points), and equipped with a linear system P' which has no base points. This gives us a map $S' \longrightarrow \mathbb{P}^N$, and its image lands in X, hence a legitimate map $S' \longrightarrow X$. □

This result gives us a way to write rational maps, and hence birational maps, as two *actual* morphisms, albeit via the inverse of one of them.

2.4 Universal property of blowing up

The other crucial fact we need characterizes the blowup.

Theorem 2.15 (Universal property of blowups): Let $f : X' \longrightarrow X$ be a birational morphism of smooth projective surfaces. Let p be a fundamental point of f^{-1}. Then f factors through the blowup $\pi : \mathrm{Bl}_p X \longrightarrow X$:

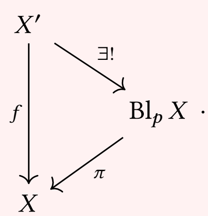

Proof. Let us call the induced birational map $T := \pi^{-1} \circ f : X' \dashrightarrow \mathrm{Bl}_p X$. If T has no fundamental points, then we are done. If T has a fundamental point q, then clearly $f(q) = p \in X$, since f is defined everywhere but π^{-1} has p as its unique fundamental point, so any fundamental point of X necessarily lies over $p \in X$. Then Proposition 2.12 implies that the total transform $\widetilde{T}(q)$ is connected and dimension ≥ 1 in $\mathrm{Bl}_p X$, while also lying over p; this means that it must be the exceptional curve $E = \pi^{-1}(p) \subset \mathrm{Bl}_p X$. So we have $\widetilde{T}(q) = E$.
Now Corollary 2.11 implies that $T^{-1} : \mathrm{Bl}_p X \dashrightarrow X'$ has finitely many fundamental points. In particular, since E is dimension 1, most points of E must be defined for T^{-1}, even though E is the total transform of a fundamental point of T! Now we can choose a point in E which is *not* a fundamental point for T^{-1} and check local coordinates; it will produce a contradiction that E is the total transform of a fundamental point of T. □

Corollary 2.16 (Birational morphisms of surfaces are a sequence of blowups): Let $f : X' \longrightarrow X$ be a birational morphism of surfaces. Let N be the number of irreducible curves in X' which are contracted to a point in X. Then N is finite, and f is the composition of exactly N blowups (at points).

Proof. The basic idea is that every time we use the universal property of blowups (2.15) to factor f into a map to a blowup, then we take care of exactly one curve, and the remaining irreducible curves are contracted in the induced map to this blowup.
If f^{-1} has no fundamental points, then f is an isomorphism, so $N = 0$ and the statement is true. Suppose f^{-1} has fundamental points; the curves which are contracted to a point must map to a fundamental point, since all non-fundamental points have at most a single point in their preimage. Then Corollary 2.11 tells us there are only finitely many fundamental points. Furthermore, the preimage of each fundamental point of f^{-1} is a closed subset of X', hence can only have finitely many irreducible components, so the set of curves which are contracted to a

point is finite.

Now let $p \in X$ be a fundamental point of f^{-1}. Then the universal property of blowing up (2.15) tells us that f factors through $\pi : \mathrm{Bl}_p X \to X$, i.e., $f = f_1 \circ \pi$ for some $f_1 : X' \to \mathrm{Bl}_p X$. Now we need to show that the number of irreducible curves contracted by f_1 is $N - 1$. First off, any curve lying over points which are not p are clearly still contracted, as π is an isomorphism away from the exceptional locus and p. Now consider the curves C lying over p. Then since $f(C) = p$, it follows that $f_1(C) \subset \pi^{-1}(p) = E$, the exceptional locus. Either $f_1(C)$ is a point, in which case, it's still contracted, or $f_1(C) = E$. We need to show that this latter case happens for exactly one curve C. But Corollary 2.11 tells us that f_1^{-1} has only finitely many fundamental points, so f_1^{-1} is defined on most of E, and in fact on a dense open subset of E, f_1^{-1} is an isomorphism. So we can conclude that there is a unique curve being mapped to E, and this curve is the closure of $f_1^{-1}(E)$ (namely, the dense open subset of E for which f_1^{-1} is defined). Therefore f_1 contracts exactly one less curve than f does. Continuing in this way, we see that f must factor into exactly N blowups. □

Remark 2.17: This is already false for smooth projective varieties X, X' of dimension ≥ 3. Let $f : X' \to X$ be the blowup of X along a smooth curve. Then every point $p \in C$ is a fundamental point of f^{-1}, since none of them have a single preimage. But f cannot possibly factor through the blowup at p, since the exceptional locus of such a blowup is dimension 2, but X' was a blowup along a curve, hence $f^{-1}(p)$ is dimension 1.

So we know that any birational *morphism* of surfaces will factor as a finite sequence of blowups. The question to be answered is what birational *maps* will look like; these are trickier, since they're not actual morphisms.

3 Main results

Here are our two main results.

3.1 Factorization of birational maps of surfaces

Theorem 3.1 (Factorization of birational maps of surfaces): Let $T : X' \dashrightarrow X$ be a birational map of surfaces. Then T can be written as the composition of blowups (at points) and their inverses. Explicitly, there exists a surface S and maps $\eta : S \to X'$, $\psi : S \to X$ such that both η, ψ are a finite sequence of blowups, and they make the following diagram commute:

$$\begin{array}{ccc} & S & \\ {}^{\eta}\swarrow & & \searrow^{\psi} \\ X' & \overset{}{\underset{T}{\dashrightarrow}} & X \end{array}.$$

Proof. First, the elimination of indeterminancy (2.14) allows us to factor T into the diagram

$$\begin{array}{ccc} & S & \\ {}^{\eta}\swarrow & & \searrow^{\psi} \\ X' & \overset{}{\underset{T}{\dashrightarrow}} & X \end{array}.$$

By construction, $\eta : S \to X'$ is already a finite sequence of blowups. It remains to see that ψ is also a sequence of blowups. But we know that $\psi : S \to X$ is a birational morphism of surfaces, so Corollary 2.16 tells us that ψ is also a sequence of blowups. □

Corollary 3.2: The arithmetic genus of a nonsingular projective surface is a birational invariant.

Proof. Arithmetic genus is unchanged by blowups at points. □

3.2 Castelnuovo's criterion

The universal property of blowups (2.15) tells us that birational morphisms factor as blowups, contracting curves to the fundamental points of the inverse map. Castelnuovo's criterion will tell us when we can realize that a surface is indeed a blowup.

Theorem 3.3 (Castelnuovo's criterion): Let Y be a curve on a (smooth projective) surface X, such that $Y \cong \mathbb{P}^1$ and $Y^2 = -1$. Then $X \cong \mathrm{Bl}_p X_0$ for some (smooth projective) surface X_0, such that Y is the exceptional curve of this blowup.

Proof. The basic idea is to construct X_0 by taking a very ample divisor H on X and modifying it slightly, so that we get a map $\varphi : X \to \mathbb{P}^n$ which is an isomorphism outside of Y, but contracts Y, thereby realizing X as the blowup of $\varphi(X)$.

In order to produce an embedding $X \hookrightarrow \mathbb{P}^N$, we need a very ample divisor H, so that the linear system $|H|$ has no base points. But we don't want a true embedding; we want it to be an embedding on $X - Y$, and then to contract Y. Therefore we modify the divisor as follows: since we want to leave $X - Y$ untouched, we take $D := H + mY$, where $m = H \cdot Y$. The reason for choosing this m will be explained shortly.

First, we need to check that this divisor D indeed defines a map $X \to \mathbb{P}^N$. To do this, we need to check that the linear system $|D|$ has no base points. First, $|D| = |H + mY| \supseteq |H| + mY$, and $|H|$ already has no base points on X, hence $|H| + mY$ has no base points on $X - Y$. This implies that $|D|$ is base-point-free away from Y. To check that it's base-point-free on Y, we want to check that $H^0(X, \mathcal{O}_X(D)) \twoheadrightarrow \mathcal{O}_X(D)|_y$ for each $y \in Y$. We can do this by first noting that $\mathcal{O}_X(D)|_Y \cong \mathcal{O}_Y(D \cdot Y) = \mathcal{O}_Y \cong \mathcal{O}_{\mathbb{P}^1}$, which is globally generated, hence $H^0(Y, \mathcal{O}_X(D)|_Y)$ surjects onto each stalk. It remains to see that $H^0(X, \mathcal{O}_X(D)) \twoheadrightarrow H^0(Y, \mathcal{O}_X(D)|_Y)$. But this can be done by noting that $\mathcal{I}_Y \cong \mathcal{O}_X(-Y)$ and using the short exact sequence

$$0 \to \mathcal{O}_X(H + (k-1)Y) \to \mathcal{O}_X(H + kY) \to \mathcal{O}_Y((H + kY) \cdot Y) \to 0,$$

applying induction, and using the long exact sequence to see that $H^1(X, \mathcal{O}_X(D - Y)) = 0$.

Now that we have a legitimate morphism $f : X \to \mathbb{P}^N$, we need to check that it has the desired geometric properties. First, by construction f is an embedding on $X - Y$. Now regarding Y, either $f(Y)$ is mapped isomorphically onto its image, or Y is contracted to a point. In the first case, $f(Y)$ would be an irreducible curve in $\mathbb{P}^N$, hence must intersect every hyperplane F, so $F \cdot f(Y) > 0$. But $f^*\mathcal{O}_{\mathbb{P}^N}(F) \cong f^*\mathcal{O}_{\mathbb{P}^N}(1) \cong \mathcal{O}_X(D)$, and $f^*F \cdot f^*f(Y) = D \cdot Y = 0$, contradiction. So f must contract Y to a point.

It remains to check that the resulting variety $f(X)$ is smooth. Clearly it's smooth outside of $f(Y)$, since f is an isomorphism on $X - Y$. So it remains to check the point $f(Y) \in f(X)$. Now we can check smoothness using the formal function theorem, but that requires $f(X)$ to be normal; so let's go ahead and take the normalization X_0 of $f(X)$, and we get a unique induced map $g : X \to X_0$, which is still an isomorphism $g : (X - Y) \xrightarrow{\sim} (X_0 - p)$, where p is the image of Y. We know that $g_*\mathcal{O}_X = \mathcal{O}_{X_0}$ from Proposition 2.5 due to g being birational and X_0 being normal. So now we can apply the formal function theorem (4.4) and check the completion of the local ring $\widehat{\mathcal{O}_{X_0,p}}$. Recall that the completion of the local ring at a point p is just a power series ring iff p is a smooth point. So we compute that

$$\widehat{\mathcal{O}_{X_0,p}} = \varprojlim_n H^0(Y_n, \mathcal{O}_{Y_n}),$$

where Y_n is the infinitesimal neighborhood of Y in X given by the ideal sheaf $\mathcal{I}_Y^n$. From here, it's a matter of computation: each $H^0(Y_n, \mathcal{O}_{Y_n}) \cong k[x, y]/(x, y)^n$, and the inverse limit of this is indeed $k[[x, y]]$, which is a power series ring, hence p is a smooth point of X_0, and hence the entire variety X_0 is nonsingular.

It remains to see that the map $g : X \to X_0$ is a blowup (or rather, blowdown). It's clear that g is a birational map which contracts exactly one curve, hence by Corollary 2.16 it must be a single blowup, namely the blowup at the point $p \in X_0$. It follows that X is realized as the blowup of a smooth projective surface, with Y as the exceptional locus. □

Remark 3.4: One can even show that $X_0 = f(X)$, i.e., that $f(X)$ was already normal, but it doesn't really matter in proving the result we want.

3.3 Examples

Example 3.5: Let $\pi : X \to C$ be a geometrically ruled surface; this just means that C is an irreducible curve, and $\pi : X \to C$ is a proper surjective map such that all of the fibers are $\mathbb{P}^1$, and there exists a section $\sigma : C \to X$. Let $p \in X$, and let $L = \pi^{-1}(\pi(p))$ be the fiber containing p. We know that any two geometrically ruled surfaces over C are birational, so Theorem 3.1 tells us that we should be able to obtain them from each other by blowups and blowdowns. We will do exactly that here: we will blow up X at p, then blow down a different exceptional curve, to obtain another geometrically ruled surface X'.
First let $f : \mathrm{Bl}_p X \to X$ be the blowup of X at point p, and let E be the exceptional locus. Let $\widetilde{L}$ be the strict transform of L. L must be birational to $\widetilde{L}$, since the blowup is an isomorphism away from p and the exceptional locus. Since $L \cong \mathbb{P}^1$ by hypothesis, then L must be birational to $\widetilde{L}$, but the only curve which is birational to $\mathbb{P}^1$ is $\mathbb{P}^1$ itself, hence $\widetilde{L} \cong \mathbb{P}^1$. We know that $L^2 = 0$, since any two fibers of π don't intersect. (More rigorously, take some function on C vanishing at $\pi(p)$; then as divisors, p is linearly equivalent to a linear combination of other points away from p; pulling back this function to X, we will see that the fiber L is linearly equivalent to a linear combination of other fibers, so L^2 is equal to the number of intersections of L with fibers over points away from p, which must be 0.) Now we see that $f^*L \sim \widetilde{L} + E$ (for example, look at set-theoretic preimage of L), so we compute that

$$0 = L^2 = (f^*L)^2 = (\widetilde{L} + E)^2 = \widetilde{L}^2 + 2\widetilde{L} \cdot E + E^2 = \widetilde{L}^2 + 2 - 1 \implies \widetilde{L}^2 = -1.$$

This means that the strict transform $\widetilde{L}$ of L is now an exceptional curve! So $\mathrm{Bl}_p X$ has two exceptional curves: the exceptional locus E, and also the strict transform $\widetilde{L}$ of the fiber L.
We are now in position to apply Castelnuovo's criterion (3.3). Since $\widetilde{L} \subset \mathrm{Bl}_p X$ is an exceptional curve, we realize $\mathrm{Bl}_p X$ as the blowup of some other surface X' at a point q, such that $\widetilde{L}$ is the exceptional locus of the blowup map $g : \mathrm{Bl}_p X \cong \mathrm{Bl}_q X' \to X'$. We just need to check that X' is indeed a geometrically ruled surface over C. It's already equipped with a map $\pi \circ f \circ g^{-1} : (X' - g(E)) \to (C - \pi(p))$, and all of these fibers are $\mathbb{P}^1$. So we just want to understand $g(E)$; this should be the last fiber over $\pi(p)$. Once again, g is an isomorphism away from $\widetilde{L}$, which intersects E at only one point, so $g(E)$ is at least birational to $E \cong \mathbb{P}^1$, hence must be $\mathbb{P}^1$. We can see that the strict transform $\widetilde{g(E)}$ of $g(E)$ is just E, pretty much by the same argument. Therefore

$$g^*(g(E)) = \widetilde{g(E)} + \widetilde{L} = E + L.$$

Squaring both sides, we find that

$$g(E)^2 = g^*(g(E))^2 = E^2 + 2E \cdot \widetilde{L} + \widetilde{L}^2 = -2 + 2 - 1 = 0.$$

So $g(E)$ is viably the fiber over $\pi(p)$ under the map we constructed $X' \dashrightarrow C$. To see that it actually *is* the fiber, we can use the valuative criterion (4.1) again. The map $X' \dashrightarrow C$ is defined on the dense open subset $X' - g(E)$, so it must be defined at the generic point. But the generic point of $g(E)$ is a codimension 1 point; the valuative criterion of properness tells us that the map $X' \dashrightarrow C$ must still be defined there, hence on all of $g(E)$, and it's clear they must all map to $\pi(p)$.

Example 3.6: Although the previous example gives us an interesting birational transformation, it is not clear that the resulting ruled surface is actually *different* from the one we started with. It turns out that sometimes we do get the same surface, but not always - we often get something new.
Let us denote by $\mathbb{F}_n$ the nth Hirzebruch surface, given by $\underline{\mathrm{Proj}}_{\mathbb{P}^1}(\mathcal{O}_{\mathbb{P}^1} \oplus \mathcal{O}_{\mathbb{P}^1}(n))$. These are ruled surfaces over $\mathbb{P}^1$, pairwise nonisomorphic for $n \geq 0$. They are characterized by the section σ having $\sigma^2 = n$. Now, the strict transform $\widetilde{\sigma} = f^*\sigma$, since it doesn't pass through p, the point being blown up. But when it's blown down to σ', the corresponding section of X', we find that $g^*\sigma' = \widetilde{\sigma} + \widetilde{L}$, from which we conclude that $(\sigma')^2 = n + 2 - 1 = n + 1 \neq (\sigma)^2$. So the birational map in the previous example transforms $\mathbb{F}_n \dashrightarrow \mathbb{F}_{n+1}$.

4 Appendix

Here, I'll list the main results we will use, but won't actually prove.

Theorem 4.1 (Valuative criterion of properness): Let $f : X \to Y$ be a noetherian morphism of finite type with X noetherian. Let K be a field. Then f is proper if and only if for every valuation ring R with fraction field K and for every morphism of $\operatorname{Spec} K$ to X and $\operatorname{Spec} R$ to Y forming the commutative diagram

$$\begin{array}{ccc} \operatorname{Spec} K & \longrightarrow & X \\ \downarrow & & \downarrow \\ \operatorname{Spec} R & \longrightarrow & Y \end{array},$$

there exists a unique map $\operatorname{Spec} R \to X$ making the diagram commute.

$$\begin{array}{ccc} \operatorname{Spec} K & \longrightarrow & X \\ \downarrow & \overset{\exists!}{\nearrow} & \downarrow \\ \operatorname{Spec} R & \longrightarrow & Y \end{array}.$$

Our main usage of the valuative criterion can be summarized as follows.

Corollary 4.2: Let $f : X \dashrightarrow Y$ be a rational map of projective varieties, proper where it is defined. Then f is defined on all codimension 1 points.

Proof. Since f is defined on a dense open subset, it must be defined at the stalk of the generic point, which is just Spec of the function field $K(X)$. But the stalk at every codimension 1 point is a discrete valuation ring in $K(X)$, so the assumption that f is proper where defined implies that f must be defined at all codimension 1 points, i.e. DVRs inside $K(X)$. □

Theorem 4.3 (Zariski main theorem): Let $f : X \to Y$ be a birational projective morphism of noetherian integral schemes with Y normal. Then every fiber of f is connected.

Theorem 4.4 (Formal function theorem): Let $f : X \to Y$ be a projective morphism of noetherian schemes, and let $\mathcal{F}$ be a coherent sheaf on X. For any $y \in Y$ and $i \geq 0$, we have an isomorphism

$$\widehat{(R^i f_* \mathcal{F})_y} \xrightarrow{\sim} \varprojlim_n H^i(X_n, \mathcal{F}_n)$$

between the completion of the local ring at y of the sheaf $R^i f_* \mathcal{F}$, and the inverse limit of the ith cohomology groups of the nth formal neighborhoods of the fiber over y. (More precisely, we define these neighborhoods to be $X_n := X \times_Y \operatorname{Spec} \mathcal{O}_{Y,y}/\mathfrak{m}_y^n$, and $\mathcal{F}_n$ to be the sheaf $\mathcal{F}$ restricted to this fiber, i.e. given by the pullback of $\mathcal{F}$ by the natural map $X_n \to X$.)

Index

References

[Bea96] Arnaud Beauville. *Complex algebraic surfaces.* Number 34. Cambridge University Press, 1996.

[Har13] Robin Hartshorne. *Algebraic geometry*, volume 52. Springer Science & Business Media, 2013.