The weight function computes the log-canonical threshold

Michael Groechenig

January 23, 2017

Contents

Introduction

These are notes for a talk given in the *Forschungsseminar* at FU Berlin, January 2017. Our main references are [NX$^+$16] and [KM08].

The goal of this talk is to explain the proof of Theorem 2.4 in [NX$^+$16]. Throughout this talk we work with an algebraically closed field k of characteristic 0. All our algebro-geometric objects will be defined relative to k, unless stated otherwise. We will denote the discretely valued field $k((t))$ by K.

In previous talks we have assigned to a regular function $Y \longrightarrow \mathbb{A}^1$ on a smooth variety Y the K-scheme $X = Y \times_{\mathbb{A}^1} \operatorname{Spec} K$. The associated Berkovich space is denoted by Y^{an}. How do the geometric structures present on Y^{an} relate to the geometry of the divisor $\Delta = f^{-1}(0)$? In the following we use the notation $\mathcal{X} = \widetilde{Y} \times_{\mathbb{A}^1} \operatorname{Spec} k[[t]]$ (where we have chosen a log-resolution $h\colon \widetilde{Y} \longrightarrow Y$ of $(Y, f^{-1}(0))$) and choose x to be a point of the divisor $f^{-1}(0)$. We denote by $\mathsf{Sk}(\mathcal{X}, x)$ the preimage $\mathsf{sp}_{\mathcal{X}}^{-1}(h^{-1}x)$. Recall that Δ gives rise to the so-called weight function $\mathsf{wt}_\Delta\colon \mathsf{Sk}(\mathcal{X}) \longrightarrow \mathbb{R}$ as defined previously in this seminar.

Theorem 0.1 (Nicaise–Xu)**.** *If wt_Δ is constant on a maximal face of the simplicial complex $\mathsf{Sk}(\mathcal{X}, x)$, then it is constant with value $\mathsf{lct}_x(Y, \Delta)$ (the log-canonical threshold of Δ at x).*

In order to understand why this is true we will do two things:

1. Recall what the log-canonical threshold is.
2. Swap the world of Berkovich spaces with the minimal model programme.

The latter is possible because we have seen that the skeleton $\mathsf{Sk}(\mathcal{X})$ admits a description in purely combinatorial terms.

1 The log-canonical threshold

Let Y be a variety and $\Delta \hookrightarrow Y$ an effective $\mathbb{Q}$-divisor. We will assume that $K_Y + \Delta$ is $\mathbb{Q}$-Cartier, that is, some integer multiple of it is a Cartier divisor.

By virtue of Hironaka's [Hir64] there exists a proper birational morphism $\widetilde{Y} \xrightarrow{h} Y$, such that $\widetilde{Y}$ is smooth, and $h^*\Delta$ is a simple normal crossing divisor. We say that h is a *log-resolution* of the pair (X, Δ). In this case we denote the h-exceptional divisors of Y by $E_1, \dots, E_k$ and write

$$h^*\Delta = \widetilde{\Delta} + \sum_{i=1}^{k} N_i E_i,$$

$$K_{\widetilde{Y}} = h^* K_X + \sum_{i=1}^{k} (\nu_i - 1) E_i.$$

Definition 1.1. *The pair (Y, Δ) is called* log-canonical*, if there exists a log-resolution $h\colon \widetilde{Y} \longrightarrow Y$, and rational numbers $a_1, \dots, a_k \geq -1$, such that*

$$K_{\widetilde{Y}} + \widetilde{\Delta} = h^*(K_Y + \Delta) + \sum_{i=1}^{k} a_i E_i.$$

A general $\mathbb{Q}$-divisor Δ is not going to be log-canonical, in fact one observes that the divisors with "large coefficients" are less likely to be log-canonical. This vague assertion can be made precise as follows: there exists a rational number $w = \mathsf{lct}_x(Y, \Delta)$, such that we have that the pair $(X, \lambda\Delta)$ is log-canonical near x if and only if $\lambda \leq w$. Let's record this as a definition.

Definition 1.2. *We define $\mathsf{lct}_x(Y, \Delta)$ to be $\sup\{\lambda \in \mathbb{Q} \,|\, (Y, \lambda\Delta)$ is log-canonical$\}$.*

A simple computation shows that a log-resolution $h\colon \widetilde{Y} \longrightarrow Y$ induces the following upper bound

$$\mathsf{lct}_x(\Delta) \leq \min\{\frac{\nu_i}{N_i} | E_i \text{ intersects } h^{-1}(x)\}.$$

It is a bit of a miracle that this a priori estimate is in fact an equality. In particular this minimum is the same for every log-resolution. A proof of this fact can be given using the $\mathbb{C}$-analogue of Igusa zeta-functions.

2 A crash course on the MMP

This section is a short introduction to the minimal model programme, told from the perspective that is relevant to us. The motivation given in the next subsection isn't actually the reason why this programme has been created in the first place. In this section we are intentionally reluctant concerning technicalities. The standing assumption is that our varieties need not be too singular, but in large places we won't specify what we actually mean by this. We only try to mentally prepare the reader for what lies ahead.

2.1 Why we need the minimal model programme

At the end of the day proving that a certain number equals the log-canonical threshold is about establishing certain inequalities for coefficients in various divisor. To us the minimal model programme is a blackbox which assigns to a projective variety Y a chain of birational maps $Y = Y_0 \dashrightarrow \cdots \dashrightarrow Y_m \dashrightarrow \cdots$, such that each birational map $Y_i \dashrightarrow Y_{i+1}$ is of a rather simple shape. The MMP is not guaranteed to terminate, if it does, then the output Y_m has one of the two properties:

(a) the canonical divisor K_{Y_m} is *nef*, that is for every curve $U \subset Y$ the intersection $U \cdot K_Y$ is non-negative,

(b) or Y_m is fibred in Fano varieties (admitting singular fibres).

Case (b) will never arise in the situations we consider in this paper. In fact, for us the main reason for running the MMP is to turn a given divisor on a variety into a nef divisor. But more on this later.

The MMP does not respect the class of smooth varieties. In order to run this "algorithm" it is necessary to permit certain types of singularieties (so-called terminal singularities). This property itself is defined in terms of inequalities for the discrepancies of a certain resolution. These inequalities are another important ingredient in the proof of Theorem 0.1.

If K_Y isn't nef, then one can find a rogue curve U, such that $U \cdot K_Y < 0$. It therefore makes sense to consider a proper morphism $Y \xrightarrow{\pi} Y_1$, such that $\pi(U)$ is a point, that is, U is contracted by π.

Definition 2.1. *Recall that $N_1(X)$ denotes the real vector space generated by integral proper curves on X, modulo the relation that $x \cdot D = y \cdot D$ for all Cartier divisors D implies $x = y \in N_1(X)$. Elements in $N_1(X)$ are also called 1-cycles. We denote by $NE(X)$ the cone of 1-cycles which have an effective representative.*

In all honesty the minimal model programme is working by contracting "extremal rays" in the cone $\overline{NE}(X)$. It turns out that these contractions $g\colon Y \longrightarrow Z$ exist in a minimal way (that is, contracting precisely R), and there are three possible cases which can arise:

(a) g is a proper birational map, such that its exceptional locus is a divisor,

(b) g is a proper birational map, such that its exceptional locus is of codimension ≥ 2 (small contraction),

(c) g is a Fano fibration.

Cases (a) and (c) are nice. The first one is also known as a *divisorial contraction*. The last one arises precisely when Y is fibred in Fano varieties. It is one of the two possible outputs of the MMP. The middle case (b) is somehow cumbersome. Due to the smallness of the exceptional locus of g, the singularities of Z are usually quite bad. It is therefore necessary to find an appropriate proper birational morphism $Z \xleftarrow{g^+} Y^+$, which is an isomorphism away from codimension 2. The resulting birational map

$$\begin{array}{ccc} Y & \dashrightarrow & Y^+ \\ & \searrow \quad \swarrow & \\ & Z & \end{array}$$

is called a *flip*. It is this flipped variety which is then fed back to the algorithm (because it's singularities are not as bad as those of Z).

Remark 2.2. *The reason for this terminology is that a flip flips the sign of the canonical divisor of* Y. *If* $g\colon Y \longrightarrow Z$ *is a small contraction then* $-K_Y$ *can be shown to be* g*-ample. If a flip exists (and in fact that's usually taken to be part of the definition), then* K_{Y^+} *is* g^+*-ample.*

Flips are the Achiles' heel of the MMP. Their existence is not known in general, and the existence of pathological examples where running the MMP leads to a never ending sequence of flips cannot be ruled out in general. Luckily the actual geometry of flips will not play a role in the proof of Theorem 0.1. We only care about the curves contracted by the map $g\colon Y \longrightarrow Z$.

2.2 Variants of the MMP

Since we aim to understand the log-canonical threshold, our main concern isn't actually the canonical divisor K_Y, but rather $K_Y + D$, where D is an effective $\mathbb{Q}$-divisor on Y. The minimal model programme is amenable to this kind of generality. A pair (Y, D) is called a *log-variety*, and we regard the sum $K_Y + D$ as *the canonical divisor of the log-variety* (Y, D). This is of course only terminology (or philosophy!).

In the log-case the minimal model programme produces a sequence of birational maps of pairs $(Y_0, D_0) \dashrightarrow \cdots \dashrightarrow (Y_m, D_m) \dashrightarrow \cdots$, and if it terminates we either have that $K_{Y_m} + D_m$ is nef, or we obtain a log-Fano fibration.

Another important version of MMP works with varieties Y which are proper over a base variety X. In this relative set-up we only work with curves (or 1-cycles) which are contracted by the structural morphism $Y \longrightarrow X$. Running the MMP leads to a sequence of proper X-varieties $Y_0, Y_1, \ldots$ and birational maps $Y_0 \dashrightarrow \cdots$ between these varieties, which are maps of X-schemes.

The proof of Theorem 0.1 requires us to combine these two varieties of MMP into one: *relative log-MMP*.

2.3 More on divisors

In this subsection we will recall the definition of various properties a pair (X, D) can satisfy. But at first we introduce a piece of terminology which will be useful now and when discussing the proof of Theorem 0.1.

Definition 2.3. *Let* $g\colon \mathbb{Q} \longrightarrow \mathbb{Q}$ *be an arbitrary function. For a* $\mathbb{Q}$*-divisor* $D = \sum_{i=1}^{k} a_i E_i$ *on a variety* X, *we write* $g(D)$ *to denote the* $\mathbb{Q}$*-divisor* $\sum_{i=1}^{k} g(a_i) E_i$.

Next we have to recall what discrepancies are. We denote by $h\colon Z \longrightarrow Y$ a proper birational morphism, where Y is normal. We assume that $K_Y + D$ is $\mathbb{Q}$-Cartier.

Definition 2.4. *For each prime divisor $E \subset Z$ of the exceptional divisor of h there exists a unique rational number $a_E(X, D)$, such that we have*

$$K_Y + \widetilde{D} = f^*(K_X + D) + \sum_E a_E(X, D)E.$$

We call the minimum of the rational numbers a_E where E runs through the prime components of the exceptional divisor of h, the discrepancy of (Y, D) along h. *We write* $\mathsf{discr}_{Y,D}(h)$.

We can now define three properties of (X, D) which are of interest to us: *klt, dlt, and lc.*

Definition 2.5. *Let (Y, D) be a log-variety, consisting of a normal variety Y and an effective $\mathbb{Q}$-divisor D.*

(a) The pair (Y, D) is said to be Kawamata log-terminal *(klt), if for every h we have* $\mathsf{discr}_{Y,D}(h) > -1$*, and* $\lfloor D \rfloor \leq 0$.

(b) The pair (Y, D) is log-canonical *(lc) if for every h as above we have* $\mathsf{discr}_{Y,D}(h) \geq -1$.

(c) If D is effective, and $0 \leq D \leq 1$, then (X, D) is said to be divisorially log-terminal *(dlt), if there exists a log-resolution h, such that* $\mathsf{discr}_{Y,D}(h) > -1$.

Our definition of dlt is somehow non-standard, and does not apply to all log-varieties (but to those we care about in this note!). See [KM08, Theorem 2.44] for why this is equivalent to the usual definition.

2.4 The Negativity Lemma

Let us recall the following well-known theorem from the theory of complex algebraic surfaces.

Lemma 2.6. *Let Y be a normal projective surface and $h\colon \widetilde{Y} \longrightarrow Y$ a resolution of singularities. We denote by $\{E_i\}_{i=1,\ldots,k} \subset \widetilde{Y}$ the exceptional divisors of this morphism (they are curves!). The matrix $(E_i \cdot E_j)_{ij}$ is* negative-definite.

A very special case is given by the blow-up of a point in $\mathbb{P}^2$. The exceptional divisor has precisely one prime component (which is a rational curve), which intersects itself with multiplicity -1.

Although this is a result in the theory of algebraic surfaces, it leads to the following higher-dimensional statement.

Lemma 2.7 (Negativity Lemma)**.** *Assume that $h\colon Z \longrightarrow Y$ is a proper birational morphism between two normal varieties. We assume that D is an h-nef divisor (that is for every curve U contracted by h, we have $U \cdot D \geq 0$).*

(a) Then $-D$ is effective if and only if $-h_*D$ is effective.

(b) If $-D$ is effective, then for every $y \in Y$ we either have that $h^{-1}(y)$ lies completely inside the support of D, or $h^{-1}(y)$ does not intersect the support of D.

We will not prove this lemma. The reader is referred to [KM08, Lemma 3.39] for more details.

3 The weight function and the log-canonical threshold

3.1 The set-up and the theorem

We can now state Theorem 0.1 using only MMP terminology. From now on we will use the notation of [NX$^+$16, Section 2].

(a) X is a smooth variety and Δ an effective $\mathbb{Q}$-divisor.

(b) $h\colon Y \longrightarrow X$ is a log-resolution of (X, Δ), such that h is an iso over $X \setminus \Delta$.

(c) For every prime component E of the exceptional divisor of h we write N_E for the E-multiplicity of $h^*\Delta$.

(c) We denote by $\nu_E - 1$ the multiplicity of $K_{Y/X} = K_Y - h^*K_X$ along E.

(d) The ratio $\frac{\nu_E}{N_E}$ is denoted by $\mathsf{wt}_\Delta(E)$. It depends only on the valuation $v_E\colon k(X) \longrightarrow \mathbb{R}$ induced by the divisor $E \subset Y$.

(e) For a finite subset $I' \subset I$ of the set of prime divisors, we write $E_{I'}$ to denote the intersection $\bigcap_{i \in I'} E_i$.

The following is Theorem 2.4 in [NX$^+$16]. We have already seen it stated in the language of Berkovich spaces in Theorem 0.1 above.

Theorem 3.1. *Let $J \subset I$ be a a non-empty subset of the set of prime divisors of $h^*\Delta$, and let C be a connected component of E_J, such that $h^{-1}(x) \cap C \neq \emptyset$, and that $J \subset I$ is maximal with respect to this property. If $\mathsf{wt}_\Delta(E_j) = \mathsf{wt}_\Delta(E_{j'}) = w$ for all $j, j' \in J$, then the value w equals the log-canonical threshold $\mathsf{lct}_x(X, \Delta)$.*

As indicated previously the proof the relative log-case of the minimal model programme. The need to work relatively arises because we have already fixed a log-resolution $h\colon Y \longrightarrow X$. The strict transform of the divisor Δ will then be appropriately modified to yield a divisor Δ_0 on Y. This modification step will be explained in detail below. It is necessary to replace $\widetilde{\Delta}$ by another divisor for a multitude of reasons. Foremost since the paper which guarantees termination of the MMP algorithm assumes certain estimates on the coefficients of the divisors which our divisor Δ wouldn't satisfy. This assumption goes by the name *klt*, that is, Kawamata log-terminal.

3.2 The proof in a nutshell

It's all about choosing the right divisor Δ_0 on Y. Our choice is constrained by the following:

(1) The relative minimal model programme for (Y, Δ_0) over X has to terminate with a model.

(2) The divisor Δ_0 should bear some relation to the strict transform $\widetilde{\Delta}$, the relative canonical divisor $K_{Y/X}$, and the number w which shows up in the statement of the theorem.

Let $\alpha\colon \mathbb{Q} \longrightarrow \mathbb{Q}$ be the function

$$\alpha(\lambda) = \begin{cases} w\lambda, & \text{if } \lambda \geq \frac{1}{w} \\ 1, & \text{otherwise.} \end{cases}$$

We define a divisor $\Delta' = \alpha(\Delta)$ on X, where we use the convention introduced in Subsection 2.3. Subsequently we set

$$\Delta_0 = \widetilde{\Delta'} + (K_{Y/X})_{\mathrm{red}}. \tag{1}$$

Why did we do this? Well, the theorem guaranteeing that relative MMP for (Y, Δ_0) terminates with a model requires a klt pair (Y, Δ_0) as input (see [BCHM10]). As we have seen, for the klt assumption to be satisfied, we need the coefficients of a divisor to be striclty < 1. So applying the function α above we get at least a non-strict inequality ≤ 1. It turns out that this is good enough for our purposes, since one can further consider a small perturbation of Δ_0 which does not affect the output of the MMP algorithm.

So in order to satisfy condition (1) above we had to work with some kind of cut-off function. In order to satisfy (2) we have included the linear part of slope w.

MMP produces a chain of birational maps $Y = Y_0 \dashrightarrow Y_1 \dashrightarrow \cdots \dashrightarrow Y_m$ over X, such that each Y_i is a $\mathbb{Q}$-factorial normal scheme over X, and with respect to the strict transform Δ_i of Δ_0 the pair (Y_i, Δ_i) is divisorially log-terminal (dlt). Furthermore, we have that $K_{Y_m} + \Delta_m$ is *nef over* X.

The key lemma which underlies the proof of 3.1 asserts that these birational maps are actually open immersions around $C \cap h^{-1}(x)$. This is only holds because of we have carefully chosen the divisor Δ_0.

Lemma 3.2 (Key Lemma)**.** *For every $0 \leq \ell \leq m$ there exists an open neighbourhood U of $C \cap h^{-1}(x)$, such that the rational map $Y_0 \dashrightarrow Y_\ell$ is an open immersion after restriction to U.*

We will return to the proof of this crucial fact in the next subsection. For now we will take it for granted and continue the proof. Let $f\colon Y_m \longrightarrow X$ denote the structural morphism of the X-scheme Y_m. By abuse of notation we write $E_j \subset Y_m$ for the image of $E_j \subset Y$ in Y_m. We begin with an easy observation.

Lemma 3.3. *The divisor $D = K_{Y_m} + \Delta_m - f^*(K_X + w\Delta)$ is nef over X.*

Proof. Recall that a divisor D on Y_m is nef over X, if for every curve $U \subset Y_m$, contracted by f, we have $U \cdot D \geq 0$. Since (Y_m, Δ_m) is the minimal model for (Y, Δ_0) produced by running the MMP, we have that $K_{Y_m} + \Delta_m$ is nef over X. We have $f^*D' \cdot U = 0$ for every divisor pulled back from X. Indeed, $f^*D' \cdot U$ is defined to be

$$\frac{\deg_U nf^*D'}{n},$$

where n is a positive integer, such that nD' is Cartier. Since the curve U is contracted by U, restricting nf^*D' to U is equal to the trivial divisor. This implies that the degree is 0. We conclude that the divisor D is nef over X. $\square$

As a next step we apply the Negativity Lemma 2.7 to show that the divisor $\Delta' = \alpha(\Delta)$ (involving the cut-off function of the linear function $\lambda \mapsto w\lambda$) is a posterio just $w\Delta$ (locally around x).

Lemma 3.4. *There exists an open neighbourhood $U \subset X$ of x, such that $D|_{f^{-1}(U)} = 0$ and $w\Delta|U = \Delta'|_U$.*

Proof. We let D denote the divisor of Lemma 3.3. We decompose Δ as $A + B$, where A and B are effective divisors, such that $wA = \alpha(A)$ and $B > \alpha(B)$. According to the Key Lemma we can choose a small open neighbourhood $V \subset Y_m$ of $f^{-1}(x) \cap C$, such that the pair $(V, \Delta_m|_U)$ can be identified an open subset of the pair (Y_0, Δ_0). Restricting to this open subset we compute

$$D = K_{Y_m} + \Delta_m - f^*(K_X + w\Delta) = K_{Y_0} + \Delta_0 - f^*(K_X + w\Delta) = D_{exc} - f_*^{-1}(wB - B_{\mathrm{red}}), \tag{2}$$

where D_{exc} denotes an f-exceptional divisor. Why? Because $K_{Y_0} - f^*K_X = K_{Y_0/X}$ is f-exeptional (since f is an isomorphism away from Δ), and $\Delta_0 - w\widetilde{\Delta} = B_{\mathrm{red}} - wB$ by definition of B.

It follows from the Negativity Lemma 2.7 that $-D$ is effective, and that if the intersection of $f^{-1}(x)$ is non-empty, then $f^{-1}(x)$ is in fact entirely contained in the support of D. Now this is the point where we use our assumption on the maximality of $J \subset I$. Indeed, this assumption implies that $f^{-1}(x) \cap C \cap E_i = \emptyset$ for $i' \in I \setminus J$. So only the divisors E_j for $j \in J$ can show up in D (locally around x). We can exclude this easily by computing the multiplicity of E_j in D. By definition it is equal to

$$\underbrace{\nu_j - 1}_{K_Y} + \underbrace{\underbrace{0}_{\widetilde{\Delta'}} + \underbrace{1}_{(K_Y)_{\mathrm{red}}}}_{\Delta_m} - \underbrace{0}_{f^*K_X} - \underbrace{wN_j}_{f^*w\Delta} = 0.$$

So in order for D to satisfy the conclusion of the Negativity Lemma we must have $D = 0$, at least over some open neighbourhood U of x. Equation (2) implies $B|U = 0$ (since $wB - B_{\mathrm{red}}$ is effective). By definition of B this shows that $w\Delta = \Delta'$ over this open neighbourhood U. □

Now we can conclude the proof of Theorem 3.1. A porism of the application of the Negativity Lemma above is that the divisors $f^*(w\Delta) - K_{Y_m} - f^*K_X$ and Δ_m are equal over an open neighbourhood of x. Since (Y_m, Δ_m) is the output of the MMP the pair (Y_m, Δ_m) is dlt. That is, choosing a log-resolution $h\colon \widetilde{Y_m} \longrightarrow Y_m$ of (Y_m, Δ_m) we only obtain discrepancies > -1. Viewing the composition $f \circ h$ as a log-resolution of (X, Δ) we therefore obtain:

$$(f \circ h)^*(K_X + w\Delta) = h^*f^*(K_X + w\Delta) = h^*\left(f^*K_X + f^*w\Delta\right) = h^*\left(K_{Y_m} - K_{Y_m/X} + f^*w\Delta\right).$$

Since $D|_U = 0$ we see that $\Delta_m = f^*w\Delta - K_{Y_m/X}$ over U. Restricting to U the right hand side is thus equivalent to

$$h^*\left(K_{Y_m} + \Delta_m\right) = K_{\widetilde{Y_m}} + \widetilde{\Delta_m} + \sum_{i \in I''} a_i E_i = K_{\widetilde{Y_m}} + w\widetilde{\Delta} + (\widetilde{\Delta_m} - w\widetilde{\Delta}) + \sum_{i \in I''} a_i E_i.$$

In order to prove that $(X, w\Delta)$ is log-canonical, it suffices to show that the discrepancy coefficients are ≥ -1. By the dlt property of (Y_m, Δ_m) the discrepancies a_i for $i \in I''$ even satisfy the strict inequality $a_i > -1$. The coefficients arising from the divisor $-\widetilde{\Delta_m}$ are ≥ -1, since all the coefficients of Δ_m are bounded by 1 by construction. Furthermore we know that 1 shows up as a coefficient of E_j for $j \in J$, which shows that for $\lambda > w$ the pair $(X, \lambda\Delta)$ can't be log-canonical in a neighbourhood of x. By the definition of the log-canonical threshold this shows that $w = \mathsf{lct}_x(X, \Delta)$.

3.3 Proof of the key lemma

We will now turn to the proof of the key lemma. That is, we have to verify that the rational maps $Y = Y_0 \dashrightarrow Y_\ell$ are actually open immersions in a sufficiently small neighbourhood of $C \cap h^{-1}(x)$ (recall that C denotes the intersection E_J). This is a case for induction. The base case $\ell = 0$ is manifestly true, since it's the identity map. Let's assume that we have established this property for a general $0 \leq \ell \leq m$, and choose a point $y \in Y_\ell$ which lies in the image of $C \cap h^{-1}(x)$ under the rational map $Y_0 \dashrightarrow Y_\ell$.

Remark 3.5. *Note that this image is not well-defined without the induction hypothesis, since a priori $Y_0 \dashrightarrow Y_\ell$ is not everywhere well-defined as a morphism. This is the reason why this argument has to be written up inductively.*

To conclude the proof we will check that $Y_\ell \dashrightarrow Y_{\ell+1}$ is an open immersion in a neighbourhood of y. Since this rational map is produced by running the MMP there exists an extremal ray R in the cone of relative curves, such that we have a strict inequality

$$R \cdot (K_{Y_\ell} + \Delta_\ell) < 0. \tag{3}$$

Let $g\colon Y_\ell \longrightarrow Z$ be the contraction of this ray R. That is, either we have $Z = Y_{\ell+1}$, or $Y_{\ell+1}$ is obtained as the flip of $Y_\ell \longrightarrow Z$.

If g contracts a curve U through y then R is generated by U. We want to show that this is not possible by exhibiting a contradiction with inequality (3).

Lemma 3.6. *If R is generated by a curve U passing through y, then the intersection number $R \cdot (K_{Y_\ell} + \Delta_\ell)$ equals* 0.

Proof. Once more we use that the intersection number of an X-relative curve with a pull-back of a divisor on X is 0 to obtain

$$R \cdot (K_{Y_\ell} + \Delta_\ell) = R \cdot (K_{Y_\ell} + \Delta_\ell - f^*(K_X + w\Delta)).$$

By virtue of the definition of Δ_ℓ as a strict transform of $\Delta_0 = \widetilde{\Delta'} + (K_{Y/X})_{\mathrm{red}}$ we can rewrite the right hand side as

$$R \cdot \left(K_{Y_\ell/X} + (K_{Y_\ell/X})_{\mathrm{red}} + \widetilde{\Delta'} - f^*w\Delta\right) = \left(\sum_{j\in J} \underbrace{(\nu_j - wN_j)}_{=0} E_j + \widetilde{B}_{\mathrm{red}} - w\widetilde{B}\right),$$

where we use again B to denote the maximal part of the effective $\mathbb{Q}$-divisor Δ, such that $wB > B$. If the intersection of R with $\widetilde{B}_{\mathrm{red}} - w\widetilde{B}$ wasn't 0, then the curve U would actually have to intersect the support of B. Arguing with so-called *log-canonical centres* one can show that this would contradict the maximality assumption of J. An introduction to log-canonical centres is given in the survey [Amb11]. We only remark that a log-canonical centre is itself a connected component of an intersection of divisors, where only those prime divisors are admitted which arise with multiplicity 1. This provides the link between these objects with the maximality assumption on J. ☐

This lemma implies that $Y_\ell \longrightarrow Z$ does not contract a curve passing through y. Therefore it follows that restricted to suitable open subsets, g is a birational proper morphism between normal varieties which is bijective. We conclude from Zariski's main theorem that g is an open immersion. Since $Y_\ell \dashrightarrow Y_{\ell+1}$ agrees with g away from the exceptional locus of g this shows that the latter map is as well an open immersion in a neighbourhood of y.

References

[Amb11] Florin Ambro, *Basic properties of log canonical centers*, Classification of algebraic varieties, 2011, pp. 39–48. 3.3

[BCHM10] Caucher Birkar, Paolo Cascini, Christopher D Hacon, and James McKernan, *Existence of minimal models for varieties of log general type*, J. Amer. Math. Soc **23** (2010), no. 2, 405–468. 3.2

[Hir64] Heisuke Hironaka, *Resolution of singularities of an algebraic variety over a field of characteristic zero: Ii*, Annals of Mathematics (1964), 205–326. 1

[KM08] János Kollár and Shigefumi Mori, *Birational geometry of algebraic varieties*, vol. 134, Cambridge university press, 2008. (document), 2.3, 2.4

[NX+16] Johannes Nicaise, Chenyang Xu, et al., *Poles of maximal order of motivic zeta functions*, Duke Mathematical Journal **165** (2016), no. 2, 217–243. (document), 3.1