Cambridge University Press
978-1-107-03534-8 - Singularities of the Minimal Model Program
János Kollár
Excerpt
More information

Introduction

In the last three decades Mori's program, the moduli theory of varieties and complex differential geometry have identified five large and important classes of singularities. These are the basic objects of this book.

Terminal. This is the smallest class needed for running Mori's program starting with smooth varieties. For surfaces, terminal equals smooth. These singularities have been fully classified in dimension 3 but they are less understood in dimensions ≥ 4.

Canonical. These are the singularities that appear on the canonical models of varieties of general type. The classification of canonical surface singularities by Du Val in 1934 is the first appearance of any of these classes in the literature. These singularities are reasonably well studied in dimension 3, less so in dimensions ≥ 4.

For many problems a modified version of Mori's program is more appropriate. Here one starts not with a variety but with a pair (X, D) consisting of a smooth variety and a simple normal crossing divisor on it. These lead to the "log" versions of the above notions.

Log terminal. This is the smallest class needed for running the minimal model program starting with a simple normal crossing pair (X, D). There are, unfortunately, many different flavors of log terminal; the above definition describes "divisorial log terminal" singularities. From the point of view of complex differential geometry, log terminal is characterized by finiteness of the volume of the smooth locus $X \setminus \operatorname{Sing} X$; that is, for any top-degree holomorphic form ω, the integral $\int_X \omega \wedge \bar{\omega}$ is finite.

Log canonical. These are the singularities that appear on the log canonical models of pairs of log general type. Original interest in these singularities came from the study of affine varieties since the log canonical model of a pair (X, D) depends only on the open variety $X \setminus D$. One can frequently view log canonical singularities as a limiting case of the log terminal ones, but they are technically much more complicated. They naturally appear in any attempt to use induction on the dimension.

The relationship of these four classes to each other seems to undergo a transition as we go from dimension 3 to higher dimensions. In dimension 3 we understand terminal singularities completely and each successive class is understood less. In dimensions ≥ 4, our knowledge about the first three classes has been about the same for a long time while very little was known about the log canonical case until recently.

Semi-log-canonical. These are the singularities that appear on the stable degenerations of smooth varieties of general type. The same way as stable degenerations of smooth curves are non-normal nodal curves, stable degenerations of higher dimensional smooth varieties also need not be normal. In essence "semi-log canonical" is the straightforward non-normal version of "log canonical," but technically they seem substantially more complicated. The main reason is that the minimal model program fails for varieties with normal crossing singularities, hence many of the basic techniques are not available.

The relationship between the study of these singularities and the development of Mori's program was rather symbiotic. Early work on the minimal models of 3-folds relied very heavily on a detailed study of 3-dimensional terminal and canonical singularities. Later developments went in the reverse direction. Several basic results, for instance adjunction theory, were first derived as consequences of the (then conjectural) minimal model program. When they were later proved independently, they provided a powerful inductive tool for the minimal model program.

Now we have relatively short direct proofs of the finite generation of the canonical rings, but several of the applications to singularity theory depend on more delicate properties of minimal models in the non-general-type case. Conversely, recent work on the abundance conjecture relies on subtle properties of semi-log canonical singularities. In writing the book, substantial effort went into untangling these interwoven threads.

The basic definitions and key results of the minimal model program are recalled in Chapter 1.

Canonical, terminal, log canonical and log terminal singularities are defined and studied in Chapter 2. As much as possible, we develop the basic theory for arbitrary schemes, rather than just for varieties over $\mathbb{C}$.

Chapter 3 contains a series of examples and classification theorems that show how complicated the various classes of singularities can be.

The technical core of the book is Chapter 4. We develop a theory of higher-codimension Poincaré residue maps and apply it to a uniform treatment of log canonical centers of arbitrary codimension. Key new innovations are the sources and springs of log canonical centers, defined in Section 4.5.

These results are applied to semi-log canonical singularities in Chapter 5. The traditional methods deal successfully with the normalization of a semi-log canonical singularity. Here we show how to descend information from the normalization of the singularity to the singularity itself.

In Chapter 6 we show that semi-log canonical singularities are Du Bois; an important property in many applications. The log canonical case was settled earlier in Kollár and Kovács (2010). With the basic properties of semi-log canonical singularities established, the induction actually runs better in the general setting.

Two properties of semi-log canonical singularities that are especially useful in moduli questions are treated in Chapter 7.

Chapter 8 is a survey of the many results about canonical, terminal, log canonical and log terminal singularities that we could not treat adequately.

Chapter 9 contains results on finite equivalence relations that were needed in previous Chapters. Some of these are technical but they should be useful in different contexts as well.

A series of auxiliary results are collected in Chapter 10.

1

Preliminaries

We usually follow the definitions and notation of Hartshorne (1977) and Kollár and Mori (1998). Those that may be less familiar or are used inconsistently in the literature are recalled in Section 1.1.

The rest of the chapter is more advanced. We suggest skipping it at first reading and then returning to these topics when they are used later.

The classical theory of minimal models is summarized in Section 1.2. Minimal and canonical models of pairs are treated in greater detail in Section 1.3. Our basic reference is Kollár and Mori (1998), but several of the results that we discuss were not yet available when Kollár and Mori (1998) appeared. In Section 1.4 we collect various theorems that can be used to improve the singularities of a variety while changing the global structure only mildly. Random facts about some singularities are collected in Section 1.5.

Assumptions Throughout this book, all schemes are assumed noetherian and separated. Further restrictions are noted at the beginning of every chapter.

All the concepts discussed were originally developed for projective varieties over $\mathbb{C}$. We made a serious effort to develop everything for rather general schemes. This has been fairly successful for the basic results in Chapter 2, but most of the later theorems are known only in characteristic 0.

1.1 Notation and conventions

Notation 1.1 The *singular locus* of a scheme X is denoted by $\operatorname{Sing} X$. It is a closed, reduced subscheme if X is excellent. The open subscheme of nonsingular points is usually denoted by X^{ns}. For regular points we use X^{reg}.

The *reduced scheme* associated to X is denoted by $\operatorname{red} X$.

Divisors and $\mathbb{Q}$-divisors

Notation 1.2 Let X be a normal scheme. A *Weil divisor*, or simply *divisor*, on X is a finite, formal, $\mathbb{Z}$-linear combination $D = \sum_i m_i D_i$ of irreducible and reduced subschemes of codimension 1. The group of Weil divisors is denoted by $\operatorname{Weil}(X)$ or by $\operatorname{Div}(X)$.

Given D and an irreducible divisor D_i, let $\operatorname{coeff}_{D_i} D$ denote the *coefficient* of D_i in D. That is, one can write $D = (\operatorname{coeff}_{D_i} D) \cdot D_i + D'$ where D_i is not a summand in D'. The *support* of D is the subscheme $\cup_i D_i \subset X$ where the union is over all those D_i such that $\operatorname{coeff}_{D_i} D \neq 0$.

A divisor D is called *reduced* if $\operatorname{coeff}_{D_i} D \in \{0, 1\}$ for every D_i. We sometimes identify a reduced divisor with its support. If $D = \sum_i a_i D_i$ (where the D_i are distinct, irreducible divisors) then $\operatorname{red} D := \sum_{i:a_i \neq 0} D_i$ denotes the reduced divisor with the same support. One can usually identify $\operatorname{red} D$ and $\operatorname{Supp} D$.

Linear equivalence of divisors is denoted by $D_1 \sim D_2$.

For a Weil divisor D, $\mathscr{O}_X(D)$ is a rank 1 reflexive sheaf and D is a Cartier divisor if and only if $\mathscr{O}_X(D)$ is locally free. The correspondence $D \mapsto \mathscr{O}_X(D)$ is an isomorphism from the group $\operatorname{Cl}(X)$ of Weil divisors modulo linear equivalence to the group of rank 1 reflexive sheaves. (This group does not seem to have a standard name but it can be identified with $\operatorname{Pic}(X \setminus \operatorname{Sing} X)$.) In this group the product of two reflexive sheaves L_1, L_2 is given by $L_1 \hat{\otimes} L_2 := \big(L_1 \otimes L_2\big)^{**}$, the double dual or reflexive hull of the usual tensor product. For powers we use the notation $L^{[m]} := (L^{\otimes m})^{**}$.

One can think of the Picard group $\operatorname{Pic}(X)$ as a subgroup of $\operatorname{Cl}(X)$.

A Weil divisor D is $\mathbb{Q}$-*Cartier* if and only if mD is Cartier for some $m \neq 0$. Equivalently, if and only if $\mathscr{O}_X(mD) = (\mathscr{O}_X(D))^{[m]}$ is locally free for some $m \neq 0$.

A normal scheme is *factorial* if every Weil divisor is Cartier and $\mathbb{Q}$-*factorial* if every Weil divisor is $\mathbb{Q}$-Cartier. See Boissière *et al.* (2011) for some foundational results.

Note further that if L is a reflexive sheaf and $D = \sum a_i D_i$ a Weil divisor then $L(D)$ denotes the sheaf of rational sections of L with poles of multiplicity at most a_i along D_i. It is thus the double dual of $L \otimes \mathscr{O}_X(D)$.

More generally, let X be a reduced, pure dimensional scheme that satisfies Serre's condition S_2. Let $\operatorname{Cl}^*(X)$ denote the abelian group generated by the irreducible Weil divisors not contained in $\operatorname{Sing} X$, modulo linear equivalence. (Thus, if X is normal, then $\operatorname{Cl}^*(X) = \operatorname{Cl}(X)$.) As before, $D \mapsto \mathscr{O}_X(D)$ is an isomorphism from $\operatorname{Cl}^*(X)$ to the group of rank 1 reflexive sheaves that are locally free at all codimension 1 points of X. For more details, see (5.6).

Aside If X is not S_2, then one should work with the group of rank 1 sheaves that are S_2. Thus $\mathscr{O}_X(\sum a_i D_i)$ should denote the sheaf of rational sections of $\mathscr{O}_X$ with poles of multiplicity at most a_i along D_i. Unfortunately, this is not consistent with the usual notation $\mathscr{O}_X(D)$ for a Cartier divisor D since on a non-S_2 scheme a locally free sheaf is not S_2, hence we will avoid using it.

Definition 1.3 ($\mathbb{Q}$-Divisors) If in the definition of a Weil divisor $\sum_i m_i D_i$ we allow $m_i \in \mathbb{Q}$ (resp. $m_i \in \mathbb{R}$), we get the notion of a $\mathbb{Q}$-*divisor* (resp. $\mathbb{R}$-*divisor*). We mostly work with $\mathbb{Q}$-divisors. For singularity theory, (2.21) reduces every question treated in this book from $\mathbb{R}$-divisors to $\mathbb{Q}$-divisors.

We say that a $\mathbb{Q}$-divisor D is a *boundary* if $0 \leq \operatorname{coeff}_{D_i} D \leq 1$ for every D_i and a *subboundary* if $\operatorname{coeff}_{D_i} D \leq 1$ for every D_i.

A $\mathbb{Q}$-divisor D is $\mathbb{Q}$-Cartier if mD is a Cartier divisor for some $m \neq 0$.

Note the difference between a $\mathbb{Q}$-Cartier divisor and a $\mathbb{Q}$-Cartier $\mathbb{Q}$-divisor.

Since the use of $\mathbb{Q}$-divisors is rather pervasive in some parts of the book, we sometimes call a divisor a $\mathbb{Z}$-*divisor* to emphasize that its coefficients are integers.

Two $\mathbb{Q}$-divisors D_1, D_2 on X are $\mathbb{Q}$-*linearly equivalent* if mD_1 and mD_2 are linearly equivalent $\mathbb{Z}$-divisors for some $m \neq 0$. This is denoted by $D_1 \sim_{\mathbb{Q}} D_2$.

Let $f \colon X \to Y$ be a morphism. Two $\mathbb{Q}$-divisors D_1, D_2 on X are *relatively* $\mathbb{Q}$-*linearly equivalent* if there is a $\mathbb{Q}$-Cartier $\mathbb{Q}$-divisor B on Y such that $D_1 \sim_{\mathbb{Q}} D_2 + f^*B$. This is denoted by $D_1 \sim_{\mathbb{Q},f} D_2$.

For a $\mathbb{Q}$-divisor $D = \sum_i a_i D_i$ (where the D_i are distinct irreducible divisors) its *round down* is $\lfloor D \rfloor := \sum_i \lfloor a_i \rfloor D_i$ where $\lfloor a \rfloor$ denotes the largest integer $\leq a$. We will also use the notation $D_{>1} =: \sum_{i:a_i>1} a_i D_i$ and similarly for $D_{<0}$, $D_{\leq 1}$ and so on.

Definition 1.4 Let $f \colon X \to S$ be a proper morphism and D a $\mathbb{Q}$-Cartier $\mathbb{Q}$-divisor on X. Let $C \subset X$ be a closed 1-dimensional subscheme of a closed fiber of f. Choose $m > 0$ such that mD is Cartier. Then

$$(D \cdot C) := \tfrac{1}{m} \deg_C(\mathscr{O}_X(mD)|_C)$$

is called the *intersection number* or *degree* of D on C.

We say that D is f-*nef* if $(D \cdot C) \geq 0$ for every such curve C. If S is the spectrum of a field, we just say that D is *nef*.

We say that D is f-*semiample* if there are proper morphisms $\pi \colon X \to Y$ and $g \colon Y \to S$ and a g-ample $\mathbb{Q}$-divisor H on Y such that $D \sim_{\mathbb{Q}} \pi^*H$. Thus f-semiample implies f-nef.

If S is a point, the difference between semiample and nef is usually minor, but for $\dim S > 0$ the distinction is frequently important; see Section 10.3.

Pairs

Mori's program was originally conceived to deal with smooth projective varieties. Later it became clear that one needs to handle certain singular varieties, schemes, algebraic or analytic spaces, and also add a divisor to the basic object.

Our main interest is in pairs (X, Δ) where

- X is either a normal variety over a field k or a normal scheme of finite type over a regular, excellent base scheme B. (In practice, the most important cases are when B is the spectrum of a field or a Dedekind ring.)
- Δ is a $\mathbb{Q}$-divisor such that $0 \leq \operatorname{coeff}_D \Delta \leq 1$ for every prime divisor D.

However, in some applications we need to work with non-normal varieties, with schemes that are not of finite type and with noneffective divisors Δ. Thus we consider the following general setting.

Definition 1.5 (Pairs) We consider pairs (X, Δ) over a base scheme B satisfying the following conditions.

(1) B is regular, excellent and pure dimensional.
(2) X is a reduced, pure dimensional, S_2, excellent scheme that has a canonical sheaf $\omega_{X/B}$ (1.6). (We will frequently simply write ω_X instead.)
(3) The canonical sheaf $\omega_{X/B}$ is locally free outside a codimension 2 subset. (This is automatic if X is normal.)
(4) $\Delta = \sum a_i D_i$ is a $\mathbb{Q}$-linear combination of distinct prime divisors none of which is contained in $\operatorname{Sing} X$. We allow the a_i to be arbitrary rational numbers. (See (2.20) for some comments on real coefficients.)

Although we will always work with schemes, the results of Chapters 1–2 all apply to algebraic spaces and to complex analytic spaces satisfying the above properties.

Comments Assumption (1) is a very mild restriction since most base schemes can be embedded into a regular scheme. However, changing the base scheme B changes $\omega_{X/B}$.

If B itself is of finite type over a field k, then we are primarily interested in the "absolute" canonical sheaf ω_X of X and not in the relative canonical sheaf $\omega_{X/B}$ for $p\colon X \to B$. There is, however, no "absolute" canonical sheaf on a scheme; the above "absolute" canonical sheaf on a k-variety is in fact $\omega_{X/\operatorname{Spec} k}$. If B is a smooth k-variety then

$$\omega_{X/B} \simeq \omega_{X/\operatorname{Spec} k} \otimes p^* \omega_B^{-1}$$

and $p^*\omega_B$ is a line bundle which is trivial along the fibers of p. In defining the singularities of Mori's program for k-varieties, we use various natural maps between various "absolute" canonical sheaves. If we work over a smooth base B, these maps just get tensored with the pull-back of ω_B^{-1}; the definitions and theorems remain unchanged. This is the reason why one can work over a regular base scheme in general.

It should be possible to define everything over a Gorenstein base scheme, but we do not see any advantage to it.

Definition 1.6 (Canonical class and canonical sheaf) For most applications, the usual definition of the canonical class and canonical sheaf given in Hartshorne (1977) and Kollár and Mori (1998) is sufficient. More generally, if X has a dualizing complex (Hartshorne (1966); Conrad (2000)), then its lowest (that is, in degree $-\dim X/S$) cohomology sheaf is the dualizing sheaf $\omega_{X/S}$. Unfortunately, the existence of a dualizing complex is a thorny problem for general schemes. Instead of getting entangled in it, we discuss a simple special case that is sufficient for most purposes. The uninitiated reader may also find the discussion in Kovács (2012b, section 5) useful.

Let B be a regular base scheme and $X \to B$ be a pure dimensional scheme of finite type over B that satisfies the following:

Condition 1.6.1 There is an open subscheme $j\colon X^0 \hookrightarrow X$ and a (locally closed) embedding $\iota\colon X^0 \hookrightarrow \mathbb{P}^N_B$ such that

(a) $Z := X \setminus X^0$ has codimension ≥ 2 in X, and
(b) $\iota(X^0)$ is a local complete intersection in $\mathbb{P}^N_B$.

Let I denote the ideal sheaf of the closure of $\iota(X^0)$. Then I/I^2 is a locally free sheaf on $\iota(X^0)$ and, as in Hartshorne (1977, II.8.20), we set

$$\omega_{X^0/B} := \iota^*(\omega_{\mathbb{P}^N/B} \otimes \det{}^{-1}(I/I^2)). \qquad (1.6.2)$$

Finally define the *canonical sheaf* of X over B as

$$\omega_{X/B} := j_*\omega_{X^0/B}. \qquad (1.6.3)$$

If $\omega_{X/B}$ is locally free, it is frequently called the *canonical bundle*. We frequently drop B from the notation.

If X^0 is smooth over B, then one can define ω_X using differentials, but in general, differential forms give a different sheaf.

(We are mainly interested in three special cases and one extension of this construction. First, if X is normal and quasi-projective then $Z = \operatorname{Sing} X$ works. Second, in dealing with stable varieties, we consider schemes X that have ordinary nodes at some codimension 1 points. Third, we occasionally use the dualizing sheaf for nonreduced divisors on a regular scheme. Finally, we sometimes use that if $p \in X$ is a point and $\hat{X}_p$ the completion of X at p then $\hat{\omega}_{X/B}$ is the canonical sheaf of $\hat{X}_p$.)

If X is reduced, the corresponding linear equivalence class of Weil divisors is denoted by K_X.

Note that while Hartshorne (1977, II.8.20) is a theorem, for us (1.6.2–3) are definitions. Therefore we need to establish that $\omega_{X/B}$ does not depend on the projective embedding chosen. This is easy to do by comparing two different embeddings ι_1, ι_2 with the diagonal embedding

$$(\iota_1, \iota_2) : X^0 \hookrightarrow \mathbb{P}^{N_1}_B \times_B \mathbb{P}^{N_2}_B \hookrightarrow \mathbb{P}^{N_1N_2+N_1+N_2}_B.$$

We also need that $\omega_{X/B}$ is the relative *dualizing sheaf.* If X itself is projective, a proof is in Kollár and Mori (1998, section 5.5). For the general case see Hartshorne (1966) and Conrad (2000).

Normal crossing conditions

Normal crossing means that something "looks like" the coordinate hyperplanes. Depending on how one interprets "looking like," one gets different notions. In many instances, the difference between them is minor or merely a technical annoyance. This is reflected by inconsistent usage in the literature. However, in some applications, especially when the base field is not algebraically closed, the differences are crucial. We have tried to adhere to the following conventions.

Definition 1.7 (Simple normal crossing for pairs) Let X be a scheme. Let $p \in X$ be a regular (not necessarily closed) point with ideal sheaf m_p and residue field $k(p)$. Then $x_1, \dots, x_n \in m_p$ are called *local coordinates* if their residue classes $\bar{x}_1, \dots, \bar{x}_n$ form a $k(p)$-basis of m_p/m_p^2.

Let $D = \sum a_i D_i$ be a Weil divisor on X. We say that (X, D) has *simple normal crossing* or *snc* at a (not necessarily closed) point $p \in X$ if X is regular at p and there is an open neighborhood $p \in X_p \subset X$ with local coordinates $x_1, \dots, x_n \in m_p$ such that $X_p \cap \operatorname{Supp} D \subset (x_1 \cdots x_n = 0)$. Alternatively, if for each D_i there is a $c(i)$ such that $D_i = (x_{c(i)} = 0)$ near p. We say that (X, D) is *simple normal crossing* or *snc* if it is snc at every point. It is important to note that being simple normal crossing is local in the Zariski topology, but not in the étale topology.

This concept is frequently called *strict normal crossing* or, if X is defined over an algebraically closed field, *global normal crossing.*

A *stratum* of an snc pair $(X, \sum_{i \in I} a_i D_i)$ is any irreducible component of an intersection $\cap_{i \in J} D_i$ for some $J \subset I$. Sometimes X itself is allowed as a stratum corresponding to $J = \emptyset$. All the strata of an snc pair are regular.

We say that (X, D) has *normal crossing* or *nc* at a point $p \in X$ if there is an étale neighborhood π: $(p' \in X') \to (p \in X)$ such that $(X', \pi^{-1} D)$ is snc at p'. Equivalently, if $(\hat{X}_K, D|_{\hat{X}_K})$ is snc at p where $\hat{X}_K$ denotes the completion at p

and K is a separable closure of $k(p)$. We say that (X, D) is *normal crossing* or *nc* if it is nc at every point. Being normal crossing is local in the étale topology.

If (X, D) is defined over a perfect field, this concept is also called *log smooth*.

Examples Let $p \in D$ be a nc point of multiplicity 2. If the characteristic is different from 2, then, in suitable local coordinates, D can be given by an equation $x_1^2 - ux_2^2 = 0$ where $u \in \mathscr{O}_{p,X}$ is a unit. D is snc at p if and only if u is a square in $\mathscr{O}_{p,X}$.

Thus $(y^2 - (1+x)x^2) \subset \mathbb{A}^2$ is nc but it is not snc at the origin. Similarly, $(x^2 + y^2 = 0) \subset \mathbb{A}^2$ is nc but it is snc only if $\sqrt{-1}$ is in the base field k.

Given (X, D), the largest open set $U \subset X$ such that $(U, D|_U)$ is snc is called the *snc locus* of (X, D). It is denoted by $\operatorname{snc}(X, D)$. Its complement is called the *non-snc locus* of (X, D) and denoted by $\operatorname{non-snc}(X, D)$.

We also use the analogously defined *nc locus,* denoted by $\operatorname{nc}(X, D)$, or its complement $\operatorname{non-nc}(X, D)$.

Finally, $p \in D$ is called a *double, triple*, etc. snc (or nc) point if D has multiplicity 2, 3, etc. at p. The *double-snc locus* (resp. the *double-nc locus*) of (X, D) is the largest open set $U \subset X$ such that $(U, D|_U)$ is snc (resp. nc) and each point of D has multiplicity ≤ 2.

Definition 1.8 (Simple normal crossing schemes) Let Y be a scheme. We say that Y has *simple normal crossing* or *snc* at a point $p \in Y$ if there is an open neighborhood $p \in Y_p \subset Y$ and a closed embedding $Y_p \hookrightarrow X_p$ of Y_p into a regular scheme X_p such that (X_p, Y_p) has simple normal crossing (1.7). We say that Y is *snc* or has *simple normal crossings* if it is snc at every point. As before, being simple normal crossing is local in the Zariski topology, but not in the étale topology.

Let X be an snc scheme with irreducible components $X = \cup_{i \in I} X_i$. A *stratum* of X is any irreducible component of an intersection $\cap_{i \in J} X_i$ for some $J \subset I$.

We say that Y has *normal crossing* or *nc* at a point $p \in Y$ if there is an étale neighborhood $\pi\colon (p' \in Y'_p) \to (p \in Y)$ and a closed embedding $Y'_p \hookrightarrow X'_p$ such that (X'_p, Y'_p) has simple normal crossing (1.7). Equivalently, if $\hat{Y}_K$ is snc at p where $\hat{Y}_K$ denotes the completion at p and K is a separable closure of $k(p)$. We say that Y is *normal crossing* or *nc* if it is nc at every point. Being normal crossing is local in the étale topology.

Note that, étale locally, snc schemes and nc schemes look the same. An nc scheme is snc if and only if its irreducible components are regular.

If X is defined over a perfect field, this concept is also called *log smooth*.