Chapter 1
Introduction

In the late 1940s Weil posed the challenge to create a cohomology theory for algebraic varieties X over an arbitrary field k with coefficients in $\mathbb{Z}$. Weil had a deep arithmetic application in mind, namely to prove a conjecture of his concerning the ζ-function attached to a variety over a finite field: Its first part asserts the rationality of this ζ-function, the second claims the existence of a functional equation, and the third predicts the locations of its zeros and poles. Assuming a well-developed cohomology theory with coefficients in $\mathbb{Z}$, possessing in particular a Lefschetz trace formula, he indicated a proof of these conjectures, cf. [48].

As we now know, for fields k of positive characteristic one cannot hope to construct meaningful cohomology groups $H^i(X, \mathbb{Z})$. Nevertheless Weil's quest provided an important stimulus. Starting with Serre and then with the main driving force of Grothendieck, several good substitutes for such a theory have been constructed. Many later developments in arithmetic algebraic geometry hinge crucially on these theories.

Among the most important cohomology theories for algebraic varieties over a field k are the following: If k is of characteristic zero, one may pass from k to the field of complex numbers $\mathbb{C}$ and consider the singular cohomology of the associated complex analytic space $X(\mathbb{C})$. Other invariants in the same situation are given by the algebraic de Rham cohomology of X. For general k and a prime ℓ different from the characteristic of k, there is the cohomology theory of étale ℓ-adic sheaves due to Grothendieck et al. All these are examples of good cohomology theories in the sense of Grothendieck, in that they possess a full set of functors f^*, $\otimes$, $f_!$, f_*, Hom, $f^!$ and a duality with certain properties. Moreover, there are important comparison isomorphisms between them.

If k is of positive characteristic p and ℓ is equal to p, one also has cohomological theories of p-adic and mod p étale sheaves. But these are not good theories in Grothendieck's sense, as they do not possess a duality and not all of the six functors above have a reasonable definition. For proper varieties over fields of characteristic p, a good theory for $\ell = p$ is crystalline cohomology.

The first significant progress towards the Weil conjecture came, however, from another approach by Dwork, who resolved the first part and some cases of the second part of the Weil conjecture by p-adic analytic methods, cf. [16]. Only later Grothendieck et al. gave a cohomological proof of this along the lines proposed by Weil. Key ingredients in this proof are the cohomological theory of ℓ-adic étale sheaves together with the Lefschetz trace formula. The latter yields an explicit formula for the L-function of an ℓ-adic étale sheaf as the L-function of a complex representing the cohomology with compact support of this sheaf. It thereby provides finer information about the L-function than Dwork's analytic proof. In 1974, Deligne gave an ingenious proof of the remaining parts of the Weil conjecture, again using extensively the cohomological method, cf. [13].

Around the same time, Drinfeld initiated an arithmetic theory for objects over global fields of positive characteristic p, cf. [14]. For this one fixes a Dedekind domain A that is finitely generated over the field with p elements $\mathbb{F}_p$ and whose group of units is finite. The simplest example is the polynomial ring $A = \mathbb{F}_q[t]$ over a finite extension with q elements $\mathbb{F}_q$. The ring A plays a role analogous to that of $\mathbb{Z}$ in classical arithmetic geometry.

The first examples of these new objects were what are now called *Drinfeld A-modules*. Related but more algebro-geometric objects, also due to Drinfeld, are *elliptic sheaves* and more generally *shtukas*. Inspired by this, in the mid 1980s Anderson introduced the notion of *t-motives* for $A = \mathbb{F}_q[t]$, which has a natural generalization to arbitrary A under the name of *A-motives*. The category of A-motives contains (via a contravariant embedding) that of Drinfeld A-modules, but is more flexible than the latter: While the only operation on Drinfeld A-modules is pullback along morphisms, on A-motives, in addition, one has all the standard operations from linear algebra such as direct sum and tensor product.

These A-motives bear many analogies to abelian varieties. For any prime ℓ of $\mathbb{Z}$ one associates the ℓ-adic Tate module to an abelian variety, and for any place v of A one associates the v-adic Tate module to an A-motive. The former is isomorphic to $\mathbb{Z}_\ell^{2g}$ if the abelian variety has dimension g, and the latter is isomorphic to A_v^r if the A-motive has rank r. In either case the Tate module carries a continuous action of the absolute Galois group of the base field. If the base field is finite, the Galois representation is completely described by the action of the Frobenius automorphism. Its main invariant is therefore the dual characteristic polynomial of Frobenius, which is an element of $1 + t\mathbb{Z}_\ell[t]$ or $1 + tA_v[t]$. One proves that this polynomial is independent of the choice of ℓ or v and lies in $1 + t\mathbb{Z}[t]$ or $1 + tA[t]$, respectively.

More generally consider a family of abelian varieties or A-motives over a base scheme X of finite type over $\mathbb{F}_p$. Following Weil, it is customary to attach an L-function to such a family by taking the product over all closed points $x \in X$ of the inverses of the above dual characteristic polynomials, stretched by the substitutions $t \mapsto t^{\deg(x)}$. This L-function is a priori a power series in $1 + t\mathbb{Z}[[t]]$ or $1 + tA[[t]]$, respectively. The first part of the Weil conjecture asserts that the L-function of a family of abelian varieties is, in fact, a rational function of t.

Inspired by the analogy with abelian varieties, Goss [22] conjectured that the L-function of a family of A-motives should be rational as well. This was proved in 1996 by Taguchi and Wan [44] for φ-sheaves, which include A-motives as a special case. Their method was inspired by Dwork's, with the field of p-adic numbers replaced by the field of Laurent series over $\mathbb{F}_p$. Taguchi and Wan also proved a large portion of a conjecture by Goss on a certain analytic L-function similar to the ζ-function of a scheme of finite type over **Spec** $\mathbb{Z}$. We shall return to this point at the end of this introduction.

With the paradigm of the development around the Weil conjecture, our main motivation for the present work was to develop a set of algebro-geometric and cohomological tools to give a purely algebraic proof of Goss's rationality conjecture. With further applications in mind that are described at the end of this introduction, we develop our theory in great generality and much detail. For the same reasons as with p-adic étale cohomology in characteristic p, we do not obtain a good cohomology theory in the sense of Grothendieck. In particular we do not obtain an $Rf^!$ or an internal Hom, and certainly there is no duality theory. But with the functors f^*, $\otimes^L$, Rf_* for proper f, and $Rf_!$ for compactifiable f, as well as the trace formula we achieve a good half, enough for the application to L-functions.

A trace formula for L-functions. One of our central results is a trace formula. Since it motivates much of the theory, we explain a basic version of it before giving more details on the individual sections of the book.

From now on we abbreviate $k := \mathbb{F}_q$, where q is a power of a prime p. Let A be a Dedekind domain that is finitely generated over k and whose group of units is finite, such as $A = k[t]$. Let X be a scheme of finite type over k and set $C :=$ **Spec** A. The Frobenius endomorphism of X relative to k, which on sections takes the form $f \mapsto f^q$, is denoted σ_X or simply σ. In what follows, schemes X, Y, etc., are thought of as base schemes for families of objects with *coefficient ring* A. We define these objects in terms of sheaves on $X \times C$, so that C plays the role of a *$\underline{C}$oefficient scheme*. As a general convention, whenever tensor or fiber products are formed over k, the subscript k will be omitted.

The basic objects of our theory are pairs $\underline{\mathcal{F}} = (\mathcal{F}, \tau)$, where $\mathcal{F}$ is a coherent sheaf on $X \times C$ and τ an $\mathcal{O}_{X\times C}$-linear homomorphism $(\sigma \times \mathrm{id})^*\mathcal{F} \to \mathcal{F}$. Such pairs are called *coherent τ-sheaves on X*. For simplicity of exposition we assume that $\mathcal{F}$ is the pullback of a coherent sheaf $\mathcal{F}_0$ on X. To any such $\underline{\mathcal{F}}$ we wish to assign an L-function as a product of pointwise L-factors.

Let $|X|$ denote the set of closed points of X. For any $x \in |X|$ let k_x denote its residue field and d_x its degree over k. Then the pullback $\mathcal{F}_x$ of $\mathcal{F}$ to $x \times C$ inherits a homomorphism $\tau_x\colon (\sigma_x \times \mathrm{id})^*\mathcal{F}_x \to \mathcal{F}_x$; hence it corresponds to a free $k_x \otimes A$-module of finite rank M_x together with a $\sigma_x \otimes \mathrm{id}_A$-linear endomorphism $\tau_x\colon M_x \to M_x$. The iterate $\tau_x^{d_x}$ of the latter is $k_x \otimes A$-linear, and one knows that

$$\det_{k_x\otimes A}(\mathrm{id} - t^{d_x}\tau_x^{d_x}|M_x) = \det_A(\mathrm{id} - t\tau_x|M_x).$$

This is therefore a polynomial in $1 + t^{d_x}A[t^{d_x}]$. Since there are at most finitely many $x \in |X|$ with fixed d_x, the following product makes sense:

Definition 1.1. The *naive L-function* of $\underline{\mathcal{F}}$ is

$$L^{\mathrm{naive}}(X, \underline{\mathcal{F}}, t) := \prod_{x\in|X|} \det_A(\mathrm{id} - t\tau_x|M_x)^{-1} \in 1 + tA[[t]].$$

For the trace formula suppose first that X is proper over k. Then for every integer i the coherent cohomology group $H^i(X, \mathcal{F}_0)$ is a finite dimensional vector space over k. Moreover, the equality $\mathcal{F} = \mathrm{pr}_1^*\mathcal{F}_0$ yields a natural isomorphism $H^i(X \times C, \mathcal{F}) \cong H^i(X, \mathcal{F}_0) \otimes A$. This is therefore a free A-module of finite rank. It also carries a natural endomorphism induced by τ; hence we can consider it as a coherent τ-sheaf on

Spec k, denoted by $H^i(X, \underline{\mathcal{F}})$. The first instance of the trace formula for L-functions then reads as follows:

Theorem 1.2. $L^{\text{naive}}(X, \underline{\mathcal{F}}, t) = \prod_{i\in\mathbb{Z}} L^{\text{naive}}(\mathbf{Spec}\, k, H^i(X, \underline{\mathcal{F}}), t)^{(-1)^i}$.

A standard procedure to extend this formula to non-proper X is via cohomology with compact support. For this we fix a dense open embedding $j : X \hookrightarrow \bar{X}$ into a proper scheme of finite type over k. We want to extend the given $\underline{\mathcal{F}}$ on X to a coherent τ-sheaf $\underline{\tilde{\mathcal{F}}}$ on $\bar{X}$ without changing the L-function. Any extension whose τ_z on $\tilde{\mathcal{F}}_z$ is zero for all $z \in |\bar{X} \setminus X|$ has that property, and it is not hard to construct one. In fact, any coherent sheaf on $\bar{X}$ extending $\mathcal{F}$, multiplied by a sufficiently high power of the ideal sheaf of $\bar{X} \setminus X$, does the job. However, there are many choices for this $\underline{\tilde{\mathcal{F}}}$, and none is functorial. Thus there is none that we can consider a natural extension by zero "$j_!\underline{\mathcal{F}}$" in the sense of τ-sheaves. Ignoring this for the moment, let us nevertheless provisionally regard $H^i(\bar{X}, \underline{\tilde{\mathcal{F}}})$ as the cohomology with compact support $H^i_c(X, \underline{\mathcal{F}})$. Then from Theorem 1.2 we obtain the following more general trace formula:

Theorem 1.3. $L^{\text{naive}}(X, \underline{\mathcal{F}}, t) = \prod_{i\in\mathbb{Z}} L^{\text{naive}}(\mathbf{Spec}\, k, H^i_c(X, \underline{\mathcal{F}}), t)^{(-1)^i}$.

Since the factors on the right hand side are polynomials in $1 + tA[t]$ or inverses of such polynomials, the *rationality* of $L^{\text{naive}}(X, \underline{\mathcal{F}}, t)$ is an immediate consequence.

We hasten to add that the order of presentation of the above theorems is for expository purposes only. We actually first prove Theorem 1.3 when X is regular and affine over k and then generalize it to arbitrary X by devissage. The proof in the affine case is based on a trace formula by Anderson from [2]. While Anderson formulated it only for $A = k$, Taguchi and Wan [45] already noted that it holds whenever A is a field, and we extend it further. Also, the formula in [2] is really the Serre dual of the one in Theorem 1.3 and therefore avoids any mention of cohomology.

The program in this book is to develop a full cohomological theory for τ-sheaves and to prove a relative trace formula for arbitrary compactifiable morphisms $f : Y \to X$. The formalism includes the inverse image functor f^*, the tensor product $\otimes$, and the direct image under a proper morphism f_*. To define an extension by zero functor $j_!$, we formally invert any morphism of coherent τ-sheaves on whose kernel and cokernel τ is nilpotent. The result is called a localization of the category of coherent τ-sheaves on X, with the same objects but other morphisms and more isomorphisms, which we call the category of *crystals* on X. In this category all the above $\underline{\tilde{\mathcal{F}}}$ become naturally isomorphic, and any one of them represents the extension by zero $j_!\underline{\mathcal{F}}$ in the sense of crystals.

This somewhat artificial construction is more than a cheap trick. It has several beneficial side effects; for example, it turns the pullback functor f^*, which is only right exact on coherent sheaves, into an exact functor on crystals. Thus crystals behave more like constructible sheaves than like coherent sheaves. The passage to crystals is perhaps the key innovation of this book compared to [2] or [45].

We construct a derived category of crystals and the corresponding derived functors f^*, $\otimes^L$ and $Rf_!$ for compactifiable f, and we calculate them in terms of the underlying τ-sheaves. We discuss flat crystals and more generally complexes of crystals of finite Tor-dimension, to which we can associate what we call crystalline L-functions. One of the main results is then the following relative trace formula generalizing Theorem 1.3.

Theorem 1.4. *Let $f : Y \to X$ be a morphism of schemes of finite type over k. Suppose that A is an integral domain that is finitely generated over k. Then for any complex $\underline{\mathcal{F}}^\bullet$ of finite Tor-dimension of crystals on Y, one has*

$$L^{\mathrm{crys}}(Y, \underline{\mathcal{F}}^\bullet, t) = L^{\mathrm{crys}}(X, Rf_!\underline{\mathcal{F}}^\bullet, t).$$

We now discuss in detail the individual chapters of the book. With the exception of Chapter 2 we follow their order of appearance. From now on $X, Y, \ldots$ will denote arbitrary noetherian schemes over k, and A will denote an arbitrary localization of a finitely generated k-algebra. In the examples A will typically be a finitely generated field, a Dedekind domain, or finite.

Basic objects (Chapter 3). The starting point is the notion of a coherent τ-sheaf, which, due to its importance, we state again:

Definition 1.5. A coherent *τ-sheaf over A on X* is a pair $\underline{\mathcal{F}} := (\mathcal{F}, \tau_{\mathcal{F}})$ consisting of a coherent sheaf $\mathcal{F}$ on $X \times C$ and an $\mathcal{O}_{X\times C}$-linear homomorphism

$$(\sigma \times \mathrm{id})^*\mathcal{F} \xrightarrow{\tau_{\mathcal{F}}} \mathcal{F}.$$

A *homomorphism* of coherent τ-sheaves $\underline{\mathcal{F}} \to \underline{\mathcal{G}}$ on X is a homomorphism of the underlying sheaves $\varphi\colon \mathcal{F} \to \mathcal{G}$ which is compatible with the action of τ.

The category of coherent τ-sheaves over A on X is denoted by $\mathbf{Coh}_\tau(X, A)$. It is an abelian A-linear category, and all constructions like kernel, cokernel, etc. are the usual ones on the underlying coherent sheaves, with the respective τ added by functoriality.

Any φ-sheaf on X in the sense of Taguchi and Wan is a coherent τ-sheaf whose underlying coherent sheaf is locally free, where A is a suitable Dedekind domain such as $k[t]$. In particular any Drinfeld A-module and any A-motive can be regarded as a coherent τ-sheaf.

The next definition is a preparation for the concept of crystals:

Definition 1.6. (a) A coherent τ-sheaf $\underline{\mathcal{F}}$ is called *nilpotent* if the iterated homomorphism $\tau^n_{\mathcal{F}}\colon (\sigma^n \times \mathrm{id})^*\mathcal{F} \to \mathcal{F}$ vanishes for some $n \gg 0$.

(b) A homomorphism of coherent τ-sheaves is called a *nil-isomorphism* if both its kernel and cokernel are nilpotent.

Nil-isomorphisms satisfy certain formal properties that make them a *multiplicative system* in the category $\mathbf{Coh}_\tau(X, A)$. To any multiplicative system one associates a certain *localized category* in which all members of the multiplicative system become isomorphisms and which has a universal property, like the localization of a ring.

Definition 1.7. The category $\mathbf{Crys}(X, A)$ of *A-crystals on X* is the localization of $\mathbf{Coh}_\tau(X, A)$ at the multiplicative system of nil-isomorphisms.

The objects of $\mathbf{Crys}(X, A)$ are those of $\mathbf{Coh}_\tau(X, A)$, but the homomorphisms are different. Homomorphisms $\underline{\mathcal{F}} \to \underline{\mathcal{G}}$ in $\mathbf{Crys}(X, A)$ are diagrams $\underline{\mathcal{F}} \leftarrow \underline{\mathcal{F}}' \to \underline{\mathcal{G}}$ in $\mathbf{Coh}_\tau(X, A)$ where the homomorphism $\underline{\mathcal{F}}' \to \underline{\mathcal{F}}$ is a nil-isomorphism, up to a certain equivalence relation. The result is again an A-linear abelian category, and a coherent τ-sheaf represents the zero crystal if and only if it is nilpotent.

Functors (Chapter 4). Consider a morphism $f : Y \to X$ and a k-algebra homomorphism $A \to A'$. Then we have the following A-linear, resp. A-bilinear, functors:

(i) inverse image: $f^*\colon \mathbf{Coh}_\tau(X, A) \to \mathbf{Coh}_\tau(Y, A)$,

(ii) tensor product: $\_ \otimes \_\colon \mathbf{Coh}_\tau(X, A) \times \mathbf{Coh}_\tau(X, A) \to \mathbf{Coh}_\tau(X, A)$,

(iii) change of coefficients: $\_ \otimes_A A'\colon \mathbf{Coh}_\tau(X, A) \to \mathbf{Coh}_\tau(X, A')$,

(iv) direct image: $f_*\colon \mathbf{Coh}_\tau(Y, A) \to \mathbf{Coh}_\tau(X, A)$ if f is proper.

All these functors are defined by the corresponding operations on the underlying coherent sheaves, with $(f \times \mathrm{id})^*$ and $(f \times \mathrm{id})_*$ in place of f^* and f_*, and with the associated τ added by functoriality. All of them preserve nil-isomorphisms and therefore pass to functors between the corresponding categories of crystals.

Next we consider an open embedding $j : U \hookrightarrow X$ with a closed complement $i : Y \hookrightarrow X$. Then we can extend any coherent τ-sheaf $\underline{\mathcal{F}}$ on U to a coherent τ-sheaf $\tilde{\underline{\mathcal{F}}}$ on X such that $i^*\tilde{\underline{\mathcal{F}}}$ is nilpotent (compare the discussion preceding Theorem 1.3). Any homomorphism between two such extensions, which is the identity on U, is a nil-isomorphism. Thus on crystals we obtain:

Theorem 1.8. *There is an exact A-linear functor*

(v) *extension by zero:* $j_!\colon \mathbf{Crys}(U, A) \to \mathbf{Crys}(X, A)$,

uniquely characterized by the properties $j^* j_! = \mathrm{id}$ *and* $i^* j_! = 0$.

Sheaf-theoretic properties (Chapter 4). Surprisingly, crystals behave more like constructible sheaves than like coherent sheaves. For instance:

Theorem 1.9. *For any morphism $f : Y \to X$ the inverse image functor f^* on crystals is exact.*

In particular, let $i_x : x \hookrightarrow X$ denote the natural embedding of a point of X. Then the *stalk at x* of a crystal $\underline{\mathcal{F}}$ on X is defined as the crystal $\underline{\mathcal{F}}_x := i_x^*\underline{\mathcal{F}}$. The following result justifies this definition:

Theorem 1.10. (a) *A sequence of crystals is exact if and only if it is exact in all stalks.*

(b) *The support of a crystal $\underline{\mathcal{F}}$, i.e., the set of points $x \in X$ for which $i_x^*\underline{\mathcal{F}}$ is non-zero* (*as a crystal!*)*, is a constructible subset of X.*

Moreover crystals enjoy a rigidity property that is not shared by τ-sheaves, but well known for étale sheaves. We state it here in its simplest form:

Theorem 1.11. *Let $f: X_{\mathrm{red}} \hookrightarrow X$ denote the canonical closed immersion of the induced reduced subscheme. Then the functors f^* and f_* are mutually quasi-inverse equivalences between* $\mathbf{Crys}(X, A)$ *and* $\mathbf{Crys}(X_{\mathrm{red}}, A)$.

Thus crystals extend in a unique way under infinitesimal extensions. In other words they grow in a prescribed way, as the name suggests.

Derived categories and functors (Chapters 5 and 6). A major part of this book deals with the extension of the functors (i)–(v) to derived functors between derived categories of crystals. As usual, the construction of derived functors requires resolutions by acyclic objects, but these cannot be found within the categories of coherent τ-sheaves or crystals. We therefore consider the much larger categories of *quasi-coherent τ-sheaves* and *quasi-crystals*, which are defined in an analogous way except that the underlying sheaves are only quasi-coherent. In these we dispose of Čech resolutions and resolutions by injectives. These categories and the functors between them are already studied in Chapters 3 and 4.

For technical reasons we also need to consider the category of all filtered direct limits of coherent τ-sheaves, which we call *ind-coherent τ-sheaves*. This category is properly sandwiched between that of coherent τ-sheaves and quasi-coherent τ-sheaves, An analogous situation occurs for crystals and quasi-crystals. A large part of Chapter 5 is devoted to clarifying the relations between the derived categories of quasi-coherent, ind-coherent and coherent τ-sheaves and of the corresponding quasi-crystals and crystals.

Once the comparison of derived categories is complete, it is relatively straightforward to construct the derived functors arising from (i)–(iv), namely derived pullback Lf^*, which equals f^* on the derived category of crystals, derived tensor product $\otimes^L$, and derived direct image under a proper morphism $Rf*$. The derived direct image with compact support $Rf_!$ for a compactifiable morphism $f = \bar{f} \circ j$ is then defined as the composite of the exact functor $j_!$ from (v) with the derived direct image under a proper morphism $R\bar{f}_*$. We establish all compatibilities between these functors that could reasonably be expected, including the proper base change theorem and the projection formula. The proofs of these formulas make essential use of the universal properties of derived functors.

We also show that the individual derived functors $L_i f^*$ and $H^{-i}(\_\otimes^L\_)$ and $R^i f_*$ for proper f can be computed in the 'naive' way by taking the corresponding derived functors of (quasi-) coherent sheaves with the respective τ added by functoriality. This is important for calculations, but would not serve as a very good definition of these functors, because one would still not know the true derived functors of crystals.

Categorical preparations (Chapter 2). We also have to deal with a number of categorical issues in relation to derived categories. Let us give three examples:

(i) There are at least two natural ways to go from the category of quasi-coherent τ-sheaves to the derived category of quasi-crystals. One can first localize at nil-isomorphisms and then derive the abelian category of quasi-crystals or, alternatively, one may first pass to the derived category of quasi-coherent τ-sheaves and then localize at nil-quasi-isomorphisms. We show that the two procedures agree.

(ii) In our constructions of derived functors, we want to stay within the setting of quasi-coherent τ-sheaves and quasi-crystals. The reader familiar with the approach in Hartshorne [25] may remember that the construction of the derived inverse image and the derived tensor product uses categories of non-quasi-coherent sheaves of modules. We bypass this by giving another criterion for the existence of derived functors.

(iii) The construction in [25] of a derived functor Rf_* of coherent sheaves for proper f depends on a comparison theorem between the derived category of coherent sheaves and the full triangulated subcategory of the derived category of quasi-coherent sheaves with coherent cohomology. The proof of this comparison theorem uses the fact that every quasi-coherent sheaf is a filtered direct limit of coherent sheaves. As mentioned above, this is not so for τ-sheaves and (quasi-) crystals. Nevertheless we prove the corresponding comparison theorem, although its proof is surprisingly intricate.

Chapter2 also contains a brief review of categorical foundations and further preparations.

Flatness (Chapter 7). Traces and characteristic polynomials and therefore L-functions can be defined only for locally free modules. The corresponding notion for crystals is flatness.

Definition 1.12. A crystal $\underline{\mathcal{F}}$ is flat if the functor $\underline{\mathcal{F}} \otimes \_$ on crystals is exact.

Clearly any crystal whose underlying coherent sheaf is locally free is flat. More importantly, any crystal on X whose underlying sheaf is a pullback of a coherent sheaf under the projection $\mathrm{pr}_1 : X \times C \to X$ is flat. Some basic properties on functorialities and on stalks are:

Theorem 1.13. (a) *Flatness of crystals is preserved under inverse image, tensor product, change of coefficients and extension by zero.*

(b) *A crystal is flat if and only if all its stalks are flat.*

(c) *If f is compactifiable and $\underline{\mathcal{F}}^\bullet$ is a bounded complex of flat crystals on Y, then $Rf_!\underline{\mathcal{F}}^\bullet$ is represented by a bounded complex of flat crystals.*

Note that by (a) the analog of assertion (c) for inverse image, tensor product and change of coefficients also holds.

We have observed that any crystal represented by a coherent τ-sheaf whose underlying sheaf is locally free is flat. For the definition of an L-function we would wish to go the other way and represent every flat crystal by such a τ-sheaf. It is a highly non-trivial matter to decide when, precisely, that is possible. As the L-function is a product of Euler factors attached to points $x \in X$ with finite residue field, it suffices to address this question over a single such point. But even then the question does not always have a positive answer. We have, however, the following important special case:

Theorem 1.14. *Suppose that $x = \mathbf{Spec}\, k_x$ for a finite field extension k_x of k and that A is artinian. Then for any A-crystal $\underline{\mathcal{F}}$ on x we have:*

(a) *The crystal* $\underline{\mathcal{F}}$ *can be represented by a coherent* τ*-sheaf* $\mathcal{F}'$ *whose homomorphism* $\tau\colon (\sigma \times \mathrm{id})^*\mathcal{F}' \to \mathcal{F}'$ *is an isomorphism.*

(b) *The representative in* (a) *is unique up to unique isomorphism. We denote it by* $\underline{\mathcal{F}}_{\mathrm{ss}}$ *and call it the* semisimple representative of $\underline{\mathcal{F}}$.

(c) *The assignment* $\underline{\mathcal{F}} \mapsto \underline{\mathcal{F}}_{\mathrm{ss}}$ *is functorial.*

(d) *The crystal* $\underline{\mathcal{F}}$ *is flat if and only if the coherent sheaf underlying* $\underline{\mathcal{F}}_{\mathrm{ss}}$ *is locally free.*

(Crystalline) L-functions (Chapters 8 and 9). In view of Theorem 1.14 (d) – and its failure for general A – we first assume that A is artinian.

Consider a flat crystal $\underline{\mathcal{F}}$ over a scheme X of finite type over k. As before let $|X|$ denote the set of closed points of X. For any $x \in |X|$ let d_x denote its degree over k and $i_x\colon x \hookrightarrow X$ its natural embedding. Then the stalk $\underline{\mathcal{F}}_x := i_x^*\underline{\mathcal{F}}$ is again flat by Theorem 1.13, and so the coherent sheaf $\mathcal{F}_{x,\mathrm{ss}}$ underlying its semisimple representative $\underline{\mathcal{F}}_{x,\mathrm{ss}}$ is locally free. Choose a locally free coherent sheaf $\mathcal{G}$ on $x \times C$ such that $\mathcal{F}_{x,\mathrm{ss}} \oplus \mathcal{G}$ is free, and turn it into a τ-sheaf $\underline{\mathcal{G}}$ by setting $\tau_{\mathcal{G}} := 0$. Then the local factor at x of the *crystalline L-function of* $\underline{\mathcal{F}}$ is defined as

$$L^{\mathrm{crys}}(x, \underline{\mathcal{F}}_x, t) := L^{\mathrm{naive}}(x, \underline{\mathcal{F}}_{x,\mathrm{ss}} \oplus \underline{\mathcal{G}}, t) \in 1 + t^{d_x} A[[t^{d_x}]]$$

with L^{naive} from Definition 1.1. One easily shows that this is independent of the choice of $\underline{\mathcal{G}}$. As there are at most finitely many $x \in |X|$ with fixed d_x, the following product makes sense:

$$L^{\mathrm{crys}}(X, \underline{\mathcal{F}}, t) := \prod_{x \in |X|} L^{\mathrm{crys}}(x, \underline{\mathcal{F}}_x, t) \in 1 + tA[[t]].$$

This definition is extended to any bounded complex $\underline{\mathcal{F}}^\bullet$ of flat A-crystals on X by defining its crystalline L-function as

$$L^{\mathrm{crys}}(X, \underline{\mathcal{F}}^\bullet, t) := \prod_{i \in \mathbb{Z}} L^{\mathrm{crys}}(x, \underline{\mathcal{F}}^i, t)^{(-1)^i} \in 1 + tA[[t]].$$

This definition is invariant under quasi-isomorphisms and thus naturally extends to any complex of crystals of bounded Tor-dimension.

Since the L-function of a crystal must be independent of the representing τ-sheaf, there is essentially no other sensible definition. In particular, before taking the characteristic polynomial one must purge any τ-subsheaf or quotient on which τ is nilpotent, because any such τ-sheaf represents the zero crystal. This has the somewhat unfortunate consequence that the naive L-function of a coherent τ-sheaf whose underlying coherent sheaf is free may differ from the crystalline L-function of the associated crystal. For example, if τ on $\mathcal{F}$ is nilpotent, the naive L-function can be any polynomial with constant term 1 and nilpotent higher coefficients, while the crystalline L-function is necessarily 1. The problem disappears only when A is reduced. It was also the reason for the name 'naive L-function'.

When A is reduced, the crystalline L-function satisfies all the usual cohomological formulas (except duality). For arbitrary artinian A, finite or not, these formulas hold only up to 'unipotent factors'. In some sense these factors correspond to missing nilpotent τ-sheaves; hence this defect is an unavoidable consequence of the very definition of crystals.

For arbitrary artinian A let $\mathfrak{n}_A$ denote its nilradical, i.e., the ideal consisting of all nilpotent elements of A. Any polynomial in $1 + t\mathfrak{n}_A[t]$ will be called *unipotent*. We regard two power series $P, Q \in 1 + tA[[t]]$ as equivalent and write $P \sim Q$ if the quotient P/Q is a unipotent polynomial. One easily checks that this defines an equivalence relation on $1 + tA[[t]]$.

Based on Theorem 1.13 (c), we can now state a main result of Chapter 9:

Theorem 1.15. *Let $f : Y \to X$ be a morphism of schemes of finite type over k. Then for any complex $\underline{\mathcal{F}}^\bullet$ of crystals on Y of bounded Tor-dimension, we have*

$$L^{\text{crys}}(Y, \underline{\mathcal{F}}^\bullet, t) \sim L^{\text{crys}}(X, Rf_!\underline{\mathcal{F}}^\bullet, t).$$

If A is reduced, then equality holds.

If we apply this to the structure morphism $X \to \mathbf{Spec}\, k$, the right hand side becomes a finite alternating product of polynomials. We therefore deduce:

Corollary 1.16. *The crystalline L-function of a complex of crystals of bounded Tor-dimension on a scheme of finite type over k is a rational function of t.*

We now revoke the assumption that A is artinian. Let Q_A denote the direct sum of the localizations of A at all minimal prime ideals. Since A is noetherian, this Q_A is artinian. Thus for any complex $\underline{\mathcal{F}}^\bullet$ of A-crystals of bounded Tor-dimension on a scheme X of finite type over k, we can consider the L-function of the associated complex of Q_A-crystals

$$L^{\text{crys}}(X, \underline{\mathcal{F}}^\bullet \otimes^L_A Q_A, t) \in 1 + tQ_A[[t]].$$

On the other hand, any good theory of L-functions should be invariant under change of coefficients. We therefore attempt to extract a crystalline L-function of $\underline{\mathcal{F}}^\bullet$ from that of $\underline{\mathcal{F}}^\bullet \otimes^L_A Q_A$. An obvious necessary condition for this is that the canonical homomorphism $A \to Q_A$ is injective. But additional conditions are needed to guarantee that the coefficients actually lie in A. For this we introduce a notion of *good coefficient rings*. Particular examples of these are normal integral domains and artinian rings. For any good coefficient ring A we show that $L^{\text{crys}}(X, \underline{\mathcal{F}}^\bullet \otimes^L_A Q_A, t)$ lies in $1 + tA[[t]]$, and so we take it as the definition of $L^{\text{crys}}(X, \underline{\mathcal{F}}^\bullet, t)$.

All formulas for crystalline L-functions, such as the trace formula in Theorem 1.15, extend directly from artinian rings to good coefficient rings. Moreover, for any ring homomorphism $\lambda \colon A \to A'$ denote the induced group homomorphism $1 + tA[[t]] \to 1 + tA'[[t]]$ again by λ. Then the change of coefficients formula for crystalline L-functions states:

Theorem 1.17. *For any homomorphism between good coefficient rings* $\lambda\colon A \to A'$ *and any complex* $\underline{\mathcal{F}}^\bullet$ *of* A*-crystals of bounded Tor-dimension on a scheme* X *of finite type over* k*, we have*

$$\lambda(L^{\mathrm{crys}}(X, \underline{\mathcal{F}}^\bullet, t)) \sim L^{\mathrm{crys}}(X, \underline{\mathcal{F}}^\bullet \otimes^L_A A', t).$$

The case of finite A**.** Our original motivation to study τ-sheaves and crystals stemmed from the theory of Drinfeld modules and A-motives. However, in relation with Artin–Schreier theory, the case of finite A was discussed in the literature already quite some time ago, cf. [29, § 4] or [SGA4$\frac{1}{2}$, § 3].

Let A be finite, and let $\underline{\mathcal{F}}$ be a coherent τ-sheaf over A on X. To its underlying coherent sheaf one can canonically assign a sheaf $\mathcal{F}_{\text{ét}}$ on the small étale site over X. By functoriality, it inherits an endomorphism τ from $\tau_{\mathcal{F}}$. The subsheaf $\varepsilon(\underline{\mathcal{F}})$ of $\mathcal{F}_{\text{ét}}$ of τ-invariant sections is an étale sheaf of A-modules, and the assignment $\underline{\mathcal{F}} \mapsto \varepsilon(\underline{\mathcal{F}})$ is functorial. Since ε preserves the property of being noetherian, it takes its image in the category $\textbf{Ét}_c(X, A)$ of constructible étale sheaves of A-modules on X. Moreover ε is left exact and zero on nilpotent τ-sheaves. Therefore it maps nil-isomorphisms to isomorphisms, and so passes to a functor

$$\bar{\varepsilon}\colon\ \textbf{Crys}(X, A) \longrightarrow \textbf{Ét}_c(X, A), \quad \underline{\mathcal{F}} \longmapsto \bar{\varepsilon}(\underline{\mathcal{F}}).$$

The following is the main result of Chapter 10. In an important special case its first part is due to Katz, cf. [29, Theorem 4.1].

Theorem 1.18. *Suppose that* A *is finite. Then the functor* $\bar{\varepsilon}$ *is an equivalence of categories. It commutes with all functors and derived functors mentioned above, and it preserves flatness. Finally, the crystalline* L*-function of a flat crystal agrees with the* L*-function of the associated étale sheaf.*

Theorem 1.18 indicates that our theory for general A may be regarded as a *global* counterpart to the theory of étale p-torsion sheaves. To be more specific, let A be a Dedekind domain that is finitely generated over k and whose group of units is finite. Then for any maximal ideal $\mathfrak{p}$ of A, the change of coefficients functor for $A \to A/\mathfrak{p}$ composed with the functor $\bar{\varepsilon}$ above induces a functor

$$\textbf{Crys}(X, A) \longrightarrow \textbf{Crys}(X, A/\mathfrak{p}) \xrightarrow{\cong} \textbf{Ét}_c(X, A/\mathfrak{p}).$$

Combining all of the above results, it follows that this functor has all compatibilities that one could hope for.

Relations to other work. For smooth schemes X, a theory similar to ours was developed independently by Emerton and Kisin in part I of [18]. The relation between the two theories is basically given by duality in the derived category of bounded complexes of coherent sheaves as described in [25]. This duality transforms $\sigma \times \mathrm{id}$-linear homomorphisms into $\sigma^{-1} \times \mathrm{id}$-linear homomorphisms and vice versa. The precise correspondence is presently worked out in detail by Blickle and the first-named author.

Emerton and Kisin go beyond our setting when, in part II of their work, they extend their theory to p-adic coefficients and prove a trace formula in that context.

The present theory is a natural extension of the theory of A-motives. Our main motivation was to prove a trace formula for L-functions of families of A-motives. The motivation of Emerton and Kisin was to prove a Riemann–Hilbert correspondence, which bears much resemblance to our last Chapter 10, and a trace formula in the p-adic context, cf. [17]. Their theory has applications to p-adic unit root crystals and local cohomology in characteristic p. Motivated by our theory, they also developed a full-fledged formalism of coefficients similar to ours and a trace formula in that context.

Further applications. The L-functions in the present book are analogues of the L-functions of p-adic and ℓ-adic sheaves on varieties in characteristic p. In [22] Goss also defined analytic L-functions of families of A-motives which are analogues of the L-functions of ℓ-adic sheaves on schemes of finite type over **Spec** $\mathbb{Z}$. He conjectured that these L-functions extend to entire resp. meromorphic functions on a domain in characteristic p which replaces the usual complex plane, and gave first evidence for this in [22]. In the case $A = \mathbb{F}_q[t]$ this conjecture was proved by Taguchi and Wan in [44] and [45]. The special values at negative integers of Goss's analytic L-functions are rational functions and turn out to have an interpretation in terms of the cohomology theory developed here. In [4] explicit cohomological expressions are used to prove that the degrees of these special values at negative integers $-n$ grow at most logarithmically in n. From this, in [4] and, independently, in [23], the remaining conjectures of Goss on meromorphy and entireness were deduced. In an interesting recent preprint [31], V. Lafforgue uses and extends parts of the present theory to study special values of Goss's L-functions at critical values.

In [5] the theory of the present book was applied to Drinfeld modular forms. The moduli scheme of Drinfeld A-modules of rank 2 with a full level $\mathfrak{n}$-structure carries a canonical locally free τ-sheaf $\underline{\mathcal{F}}_{\mathfrak{n}}$ of rank 2, which corresponds to the universal Drinfeld module. Following the classical case of elliptic modular forms, one is naturally led to studying the cohomology of the k-th symmetric power $\mathrm{Sym}^k\, \underline{\mathcal{F}}_{\mathfrak{n}}$ for any $k \geq 0$. In [5] the 'analytic realization' of this cohomology is shown to be dual to the space of cuspidal Drinfeld modular forms of weight $k+2$ and level $\mathfrak{n}$. This isomorphism is equivariant for naturally defined Hecke operators. On the other hand, the 'étale realization' of this cohomology yields abelian A_v-adic Galois representation for any place v of A. Combining both realizations, any cuspidal Drinfeld Hecke eigenform yields a one-dimensional v-adic Galois representation for any v, where the correspondence is given by an Eichler–Shimura type relation. This is analogous to constructions by Eichler–Shimura and Deligne attaching ℓ-adic Galois representation to cuspidal elliptic modular forms. For a survey on both developments we refer to [6].