Chapter 2
Ultraproducts and Łoś' Theorem

In this chapter, W denotes an infinite set, always used as an index set, on which we fix a non-principal ultrafilter.[1] Given any collection of (first-order) structures indexed by W, we can define their ultraproduct. However, in this book, we will be mainly concerned with the construction of an ultraproduct of rings, an *ultra-ring* for short, which is then defined as a certain residue ring of their Cartesian product. From this point of view, the construction is purely algebraic, although it is originally a model-theoretic one (we only provide some supplementary background on the model-theoretic perspective). We review some basic properties (deeper theorems will be proved in the later chapters), the most important of which is Łoś' Theorem, relating properties of the approximations with their ultraproduct. When applied to algebraically closed fields, we arrive at a result that is pivotal in most of our applications: the Lefschetz Principle (Theorem 2.4.3), allowing us to transfer many properties between positive and zero characteristic.

2.1 Ultraproducts

We start with the classical definition of ultraproducts via ultrafilters; for different approaches, see §§2.5 and 2.6 below.

2.1.1 Ultrafilters

By a (*non-principal*) *ultrafilter* $\mathfrak{W}$ on W, we mean a collection of infinite subsets of W closed under finite intersection, with the property that for any subset $D \subseteq W$, either D or its complement $-D$ belongs to $\mathfrak{W}$. In particular, the empty set does not belong to $\mathfrak{W}$, and if $D \in \mathfrak{W}$ and E is an arbitrary set containing D, then also

[1] We will drop the adjective 'non-principal' since these are the only ultrafilters we are interested in; if we want to talk about principal ones, we just say *principal filter*; and if we want to talk about both, we say *maximal filter*.

H. Schoutens, *The Use of Ultraproducts in Commutative Algebra*,
Lecture Notes in Mathematics 1999, DOI 10.1007/978-3-642-13368-8_2,

$E \in \mathfrak{W}$, for otherwise $-E \in \mathfrak{W}$, whence $\emptyset = D \cap -E \in \mathfrak{W}$, contradiction. Since every set in $\mathfrak{W}$ must be infinite, it follows that any co-finite set belongs to $\mathfrak{W}$. The existence of ultrafilters follows from the Axiom of Choice, and we make this set-theoretic assumption henceforth. It follows that for any infinite subset of W, we can find an ultrafilter containing this set.

More generally, a proper collection of subsets of W is called a *filter* if it is closed under intersection and supersets. Any ultrafilter is a filter which is maximal with respect to inclusion. If we drop the requirement that all sets in $\mathfrak{W}$ must infinite, then some singleton must belong to $\mathfrak{W}$; such a filter is called *principal*, and these are the only other maximal filters. A maximal filter is an ultrafilter if and only if it contains the *Frechet filter* consisting of all co-finite subsets (for all these properties, see for instance [81, §4] or [57, §6.4]).

In the remainder of these notes, unless stated otherwise, we fix an ultrafilter $\mathfrak{W}$ on W, and (almost always) omit reference to this fixed ultrafilter from our notation. No extra property of the ultrafilter is assumed, with the one exception described in Remark 8.1.5, which is nowhere used in the rest of our work anyway. Ultrafilters play the role of a decision procedure on the collection of subsets of W by declaring some subsets 'large' (those belonging to $\mathfrak{W}$) and declaring the remaining ones 'small'. More precisely, let o_w be elements indexed by $w \in W$, and let $\mathscr{P}$ be a property. We will use the expressions *almost all o_w satisfy property $\mathscr{P}$* or *o_w satisfies property $\mathscr{P}$ for almost all w* as an abbreviation of the statement that there exists a set D in the ultrafilter $\mathfrak{W}$, such that property $\mathscr{P}$ holds for the element o_w, whenever $w \in D$. Note that this is also equivalent with the statement that the set of all $w \in W$ for which o_w has property $\mathscr{P}$, lies in the ultrafilter (read: *is large*).

2.1.2 Ultraproducts

Let O_w be sets, for $w \in W$. We define an equivalence relation on the Cartesian product $O_\infty := \prod O_w$, by calling two sequences (a_w) and (b_w), for $w \in W$, equivalent, if a_w and b_w are equal for almost all w. In other words, if the set of indices $w \in W$ for which $a_w = b_w$ belongs to the ultrafilter. We will denote the equivalence class of a sequence (a_w) by

$$\operatorname*{ulim}_{w\to\infty} a_w, \quad \text{or} \quad \operatorname{ulim} a_w, \quad \text{or} \quad a_\natural.$$

The set of all equivalence classes on $\prod O_w$ is called the *ultraproduct* of the O_w and is denoted

$$\operatorname*{ulim}_{w\to\infty} O_w, \quad \text{or} \quad \operatorname{ulim} O_w, \quad \text{or} \quad O_\natural.$$

If all O_w are equal to the same set O, then we call their ultraproduct the *ultrapower* $O_\natural$ of O. There is a canonical map $O \to O_\natural$, sometimes called the *diagonal embedding*, sending an element o to the image of the constant sequence o in $O_\natural$. To see that it is an injection, assume o' has the same image as o in $O_\natural$. This means that for almost all w, and hence for at least one, the elements o and o' are equal.

Note that the element-wise and set-wise notations are reconciled by the fact that

$$\operatorname*{ulim}_{w\to\infty}\{o_w\} = \{\operatorname*{ulim}_{w\to\infty} o_w\}.$$

The more common notation for an ultraproduct one usually finds in the literature is O^*; in the past, I also have used O_∞, which in this book is reserved to denote Cartesian products. The reason for using the particular notation $O_\natural$ in these notes is because we will also introduce the remaining chromatic products $O_\flat$ and $O_\sharp$ (at least for certain local rings; see Chapters 9 and 8 respectively).

We will also often use the following terminology: if o is an element in an ultraproduct $O_\natural$, then any choice of elements $o_w \in O_w$ with ultraproduct equal to o will be called an *approximation* of o. Although an approximation is not uniquely determined by the element, any two agree almost everywhere. Below we will extend our usage of the term approximation to include other objects as well.

2.1.3 Properties of Ultraproducts

For the following properties, the easy proofs of which are left as an exercise, let O_w be sets with ultraproduct $O_\natural$.

2.1.1 *If Q_w is a subset of O_w for each w, then* $\operatorname{ulim} Q_w$ *is a subset of $O_\natural$.*

In fact, $\operatorname{ulim} Q_w$ consists of all elements of the form $\operatorname{ulim} o_w$, with almost all o_w in Q_w.

2.1.2 *If each O_w is the graph of a function $f_w\colon A_w \to B_w$, then $O_\natural$ is the graph of a function $A_\natural \to B_\natural$, where $A_\natural$ and $B_\natural$ are the respective ultraproducts of A_w and B_w. We will denote this function by*

$$\operatorname*{ulim}_{w\to\infty} f_w \quad \textit{or} \quad f_\natural.$$

Moreover, we have an equality

$$\operatorname*{ulim}_{w\to\infty}(f_w(a_w)) = (\operatorname*{ulim}_{w\to\infty} f_w)(\operatorname*{ulim}_{w\to\infty} a_w), \tag{2.1}$$

for $a_w \in A_w$.

2.1.3 *If each O_w comes with an operation $*_w\colon O_w \times O_w \to O_w$, then*

$$*_\natural := \operatorname*{ulim}_{w\to\infty} *_w$$

is an operation on $O_\natural$. If all (or, almost all) O_w are groups with multiplication $*_w$ and unit element 1_w, then $O_\natural$ is a group with multiplication $*_\natural$ and unit element $1_\natural :=$ $\operatorname{ulim} 1_w$. If almost all O_w are Abelian groups, then so is $O_\natural$.

2.1.4 *If each O_w is a (commutative) ring under the addition $+_w$ and the multiplication $\cdot_w$, then $O_\natural$ is a (commutative) ring with addition $+_\natural$ and multiplication $\cdot_\natural$.*

In fact, in that case, $O_\natural$ is just the quotient of the product $O_\infty := \prod O_w$ modulo the *null-ideal*, the ideal consisting of all sequences (o_w) for which almost all o_w are zero (for more on this ideal, see §2.5 below). From now on, we will drop subscripts on the operations and denote the ring operations on the O_w and on $O_\natural$ simply by $+$ and $\cdot$.

2.1.5 *If almost all O_w are fields, then so is $O_\natural$. More generally, if almost each O_w is a domain with field of fractions K_w, then the ultraproduct $K_\natural$ of the K_w is the field of fractions of $O_\natural$.*

Just to give an example of how to work with ultraproducts, let me give the proof: if $a \in O_\natural$ is non-zero, with approximation a_w (recall that this means that $\operatorname{ulim} a_w = a$), then by the previous description of the ring structure on $O_\natural$, almost all a_w will be non-zero. Therefore, letting b_w be the inverse of a_w whenever this makes sense, and zero otherwise, one verifies that $\operatorname{ulim} b_w$ is the inverse of a. □

2.1.6 *If C_w are rings and O_w is an ideal in C_w, then $O_\natural$ is an ideal in $C_\natural := \operatorname{ulim} C_w$. In fact, $O_\natural$ is equal to the subset of all elements of the form $\operatorname{ulim} o_w$ with almost all $o_w \in O_w$. Moreover, the ultraproduct of the C_w/O_w is isomorphic to $C_\natural/O_\natural$. If almost every O_w is generated by e elements, then so is $O_\natural$.*

In other words, the ultraproduct of ideals $O_w \subseteq C_w$ is equal to the image of the ideal $\prod O_w$ in the product $C_\infty := \prod C_w$ under the canonical residue homomorphism $C_\infty \to C_\natural$. As for the last assertion, suppose $o_{1w}, \dots, o_{e(w),w}$ generate O_w, for each w, and let $o_{i\natural}$ be the ultraproduct of the $o_{i,w}$, where we put the latter equal to 0 if $i > e(w)$. The ideal generated by the $o_{i\natural}$ can be strictly contained in $O_\natural$ (an example is the ideal of infinitesimals, defined below in 2.4.13), but it is equal to it if almost all e_w are equal, say, to e. Indeed, any element $o_\natural$ of $O_\natural$ is an ultraproduct of elements $o_w \in O_w$, which therefore can be written as a linear combination $o_w = r_{1w}o_{1w} + \dots + r_{e,w}o_{e,w}$, for some $r_{i,w} \in C_w$. Let $r_{i\natural} \in C_\natural$ be the ultraproduct of the $r_{i,w}$, for $i = 1, \dots, e$. By Łoś' Theorem (see Theorem 2.3.2 below), we have $o_\natural = r_{1\natural}o_{1\natural} + \dots + r_{e\natural}o_{e\natural}$.

2.1.7 *If $f_w\colon A_w \to B_w$ are ring homomorphisms, then the ultraproduct $f_\natural$ is again a ring homomorphism. In particular, if σ_w is an endomorphism on A_w, then the ultraproduct $\sigma_\natural$ is a ring endomorphism on $A_\natural := \operatorname{ulim} A_w$.*

2.2 Model-theory in Rings

The previous examples are just instances of the general principle that 'algebraic structure' carries over to the ultraproduct. The precise formulation of this principle is called *Łoś' Theorem* (Łoś is pronounced 'wôsh') and requires some terminology from model-theory. However, for our purposes, a weak version of Łoś' Theorem (namely Theorem 2.3.1 below) suffices in almost all cases, and its proof is entirely algebraic. Nonetheless, for a better understanding, the reader is invited to indulge in some elementary model-theory, or rather, an ad hoc version for rings only (if this not satisfies him/her, (s)he should consult any textbook, such as [57, 67, 81]).

2.2.1 Formulae

By a *quantifier free formula without parameters* in the free variables $\xi = (\xi_1, \dots, \xi_n)$, we will mean an expression of the form

$$\varphi(\xi) := \bigvee_{j=1}^{m} f_{1j} = 0 \wedge \ldots \wedge f_{sj} = 0 \wedge g_{1j} \neq 0 \wedge \ldots \wedge g_{tj} \neq 0, \tag{2.2}$$

where each f_{ij} and g_{ij} is a polynomial with integer coefficients in the variables ξ, and where $\wedge$ and $\vee$ are the logical connectives *and* and *or*. If instead we allow the f_{ij} and g_{ij} to have coefficients in a ring R, then we call $\varphi(\xi)$ a *quantifier free formula with parameters in R*. We allow all possible degenerate cases as well: there might be no variables at all (so that the formula simply declares that certain elements in $\mathbb{Z}$ or in R are zero and others are non-zero) or there might be no equations or no negations or perhaps no conditions at all. Put succinctly, a quantifier free formula is a Boolean combination of polynomial equations using the connectives $\wedge$, $\vee$ and $\neg$ (negation), with the understanding that we use distributivity and De Morgan's Laws to rewrite this Boolean expression in the (disjunctive normal) form (2.2).

By a *formula without parameters* in the free variables ξ, we mean an expression of the form

$$\varphi(\xi) := (\mathrm{Q}_1\, \zeta_1) \cdots (\mathrm{Q}_p\, \zeta_p)\, \psi(\xi, \zeta),$$

where $\psi(\xi, \zeta)$ is a quantifier free formula without parameters in the free variables ξ and $\zeta = (\zeta_1, \dots, \zeta_p)$ and where Q_i is either the universal quantifier $\forall$ or the existential quantifier $\exists$. If instead $\psi(\xi, \zeta)$ has parameters from R, then we call $\varphi(\xi)$ a *formula with parameters* in R. A formula with no free variables is called a *sentence*.

2.2.2 Satisfaction

Let $\varphi(\xi)$ be a formula in the free variables $\xi = (\xi_1, \dots, \xi_n)$ with parameters from R (this includes the case that there are no parameters by taking $R = \mathbb{Z}$ and the case that there are no free variables by taking $n = 0$). Let A be an R-algebra and let

$\mathbf{a} = (a_1, \dots, a_n)$ be a tuple with entries from A. We will give meaning to the expression $\mathbf{a}$ *satisfies* the formula $\varphi(\xi)$ in A (sometimes abbreviated to $\varphi(\mathbf{a})$ *holds in* A or *is true in* A) by induction on the number of quantifiers. Suppose first that $\varphi(\xi)$ is quantifier free, given by the Boolean expression (2.2). Then $\varphi(\mathbf{a})$ holds in A, if for some j_0, all $f_{ij_0}(\mathbf{a}) = 0$ and all $g_{ij_0}(\mathbf{a}) \neq 0$. For the general case, suppose $\varphi(\xi)$ is of the form $(\exists\zeta)\,\psi(\xi,\zeta)$ (respectively, $(\forall\zeta)\,\psi(\xi,\zeta)$), where the satisfaction relation is already defined for the formula $\psi(\xi,\zeta)$. Then $\varphi(\mathbf{a})$ holds in A, if there is some $b \in A$ such that $\psi(\mathbf{a},b)$ holds in A (respectively, if $\psi(\mathbf{a},b)$ holds in A, for all $b \in A$). The subset of A^n consisting of all tuples satisfying $\varphi(\xi)$ will be called the *subset defined* by φ, and will be denoted $\varphi(A)$. Any subset that arises in such way will be called a *definable subset* of A^n.

Note that if $n = 0$, then there is no mention of tuples in A. In other words, a sentence is either true or false in A. By convention, we set A^0 equal to the singleton $\{\emptyset\}$ (that is to say, A^0 consists of the empty tuple $\emptyset$). If φ is a sentence, then the set defined by it is either $\{\emptyset\}$ or $\emptyset$, according to whether φ is true or false in A.

2.2.3 Constructible Sets

There is a connection between definable subsets and Zariski-constructible subsets, where the relationship is the most transparent over algebraically closed fields, as we will explain below. In general, we can make the following observations.

Let R be a ring. Let $\varphi(\xi)$ be a quantifier free formula with parameters from R, given as in (2.2). Let $\Sigma_{\varphi(\xi)}$ denote the constructible subset of $\mathbb{A}^n_R = \operatorname{Spec}(R[\xi])$ consisting of all prime ideals $\mathfrak{p}$ which, for some j_0, contain all f_{ij_0} and do not contain any g_{ij_0}. In particular, if $n = 0$, so that $\mathbb{A}^0_R$ is by definition $\operatorname{Spec}(R)$, then the constructible subset Σ_φ associated to φ is a subset of $\operatorname{Spec}(R)$.

Let A be an R-algebra and assume moreover that A is a domain (we will never use constructible sets associated to formulae if A is not a domain). For an n-tuple $\mathbf{a}$ over A, let $\mathfrak{p}_\mathbf{a}$ be the (prime) ideal in $A[\xi]$ generated by the $\xi_i - a_i$, where $\xi = (\xi_1, \dots, \xi_n)$. Since $A[\xi]/\mathfrak{p}_\mathbf{a} \cong A$, we call such a prime ideal an *A-rational point* of $A[\xi]$. It is not hard to see that this yields a bijection between n-tuples over A and A-rational points of $A[\xi]$, which we therefore will identify with one another. In this terminology, $\varphi(\mathbf{a})$ holds in A if and only if the corresponding A-rational point $\mathfrak{p}_\mathbf{a}$ lies in the constructible subset $\Sigma_{\varphi(\xi)}$ (strictly speaking, we should say that it lies in the base change $\Sigma_{\varphi(\xi)} \times_{\operatorname{Spec}(R)} \operatorname{Spec}(A)$, but for notational clarity, we will omit any reference to base changes). If we denote the collection of A-rational points of the constructible set $\Sigma_{\varphi(\xi)}$ by $\Sigma_{\varphi(\xi)}(A)$, then this latter set corresponds to the definable subset $\varphi(A)$ under the identification of A-rational points of $A[\xi]$ with n-tuples over A. If φ is a sentence, then Σ_φ is a constructible subset of $\operatorname{Spec}(R)$ and hence its base change to $\operatorname{Spec}(A)$ is a constructible subset of $\operatorname{Spec}(A)$. Since A is a domain, $\operatorname{Spec}(A)$ has a unique A-rational point (corresponding to the zero-ideal) and hence φ holds in A if and only if this point belongs to Σ_φ.

Conversely, if Σ is an R-constructible subset of $\mathbb{A}^n_R$, then we can associate to it a quantifier free formula $\varphi_\Sigma(\xi)$ with parameters from R as follows. However, here there is some ambiguity, as a constructible subset is more intrinsically defined than a formula. Suppose first that Σ is the Zariski closed subset $\mathrm{V}(I)$, where I is an ideal in $R[\xi]$. Choose a system of generators, so that $I = (f_1, \dots, f_s)R[\xi]$ and set $\varphi_\Sigma(\xi)$ equal to the quantifier free formula $f_1(\xi) = \dots = f_s(\xi) = 0$. Let A be an R-algebra without zero-divisors. It follows that an n-tuple $\mathbf{a}$ is an A-rational point of Σ if and only if $\mathbf{a}$ satisfies the formula φ_Σ. Therefore, if we make a different choice of generators $I = (f'_1, \dots, f'_s)R[\xi]$, although we get a different formula φ', it defines in any R-algebra A without zero-divisors the same definable subset, to wit, the collection of A-rational points of Σ. To associate a formula to an arbitrary constructible subset, we do this recursively by letting $\varphi_\Sigma \wedge \varphi_\Psi$, $\varphi_\Sigma \vee \varphi_\Psi$ and $\neg\varphi_\Sigma$ correspond to the constructible sets $\Sigma \cap \Psi$, $\Sigma \cup \Psi$ and $-\Sigma$ respectively.

We say that two formulae $\varphi(\xi)$ and $\psi(\xi)$ in the same free variables $\xi = (\xi_1, \dots, \xi_n)$ are *equivalent* over a ring A, if they hold on exactly the same tuples from A (that is to say, if they define the same subsets in A^n). In particular, if φ and ψ are sentences, then they are equivalent in A if they are simultaneously true or false in A. If $\varphi(\xi)$ and $\psi(\xi)$ are equivalent for all rings A in a certain class $\mathscr{K}$, then we say that $\varphi(\xi)$ and $\psi(\xi)$ are *equivalent modulo the class* $\mathscr{K}$. In particular, if Σ is a constructible subset in $\mathbb{A}^n_R$, then any two formulae associated to it are equivalent modulo the class of all R-algebras without zero-divisors. In this sense, there is a one-one correspondence between constructible subsets of $\mathbb{A}^n_R$ and quantifier free formulae with parameters from R up to equivalence.

2.2.4 Quantifier Elimination

For certain rings (or classes of rings), every formula is equivalent to a quantifier free formula; this phenomenon is known under the name *Quantifier Elimination*. We will only encounter it for the following class.

Theorem 2.2.1 (Quantifier Elimination for Algebraically Closed Fields). *If $\mathscr{K}$ is the class of all algebraically closed fields, then any formula without parameters is equivalent modulo $\mathscr{K}$ to a quantifier free formula without parameters.*

More generally, if F is a field and $\mathscr{K}(F)$ the class of all algebraically closed fields containing F, then any formula with parameters from F is equivalent modulo $\mathscr{K}(F)$ to a quantifier free formula with parameters from F.

Proof (Sketch of proof). These statements can be seen as translations in model-theoretic terms of Chevalley's Theorem which says that the projection of a constructible subset is again constructible. I will only explain this for the first assertion. As already observed, a quantifier free formula $\varphi(\xi)$ (without parameters) corresponds to a constructible set $\Sigma_{\varphi(\xi)}$ in $\mathbb{A}^n_\mathbb{Z}$ and the tuples in K^n satisfying $\varphi(\xi)$ are precisely the K-rational points $\Sigma_{\varphi(\xi)}(K)$ of $\Sigma_{\varphi(\xi)}$. The key observation is now the following. Let $\psi(\xi, \zeta)$ be a quantifier free formula and put

$\gamma(\xi) := (\exists\zeta)\,\psi(\xi,\zeta)$, where $\xi = (\xi_1,\dots,\xi_n)$ and $\zeta = (\zeta_1,\dots,\zeta_m)$. Let $\Psi := \psi(K)$ be the subset of K^{n+m} defined by $\psi(\xi,\zeta)$ and let $\Gamma := \gamma(K)$ be the subset of K^n defined by $\gamma(\xi)$. Therefore, if we identify K^{n+m} with the collection of K-rational points of $\mathbb{A}_K^{n+m}$, then

$$\Psi = \Sigma_{\psi(\xi,\zeta)}(K).$$

Moreover, if $p\colon \mathbb{A}_K^{n+m} \to \mathbb{A}_K^n$ is the projection onto the first n coordinates then $p(\Psi) = \Gamma$. By Chevalley's Theorem (see for instance [27, Corollary 14.7] or [39, II. Exercise 3.19]), $p(\Sigma_{\psi(\xi,\zeta)})$ (as a subset in $\mathbb{A}_{\mathbb{Z}}^n$) is again constructible, and therefore, by our previous discussion, of the form $\Sigma_{\chi(\xi)}$ for some quantifier free formula $\chi(\xi)$. Hence $\Gamma = \Sigma_{\chi(\xi)}(K)$, showing that $\gamma(\xi)$ is equivalent modulo K to $\chi(\xi)$. Since $\chi(\xi)$ does not depend on K, we have in fact an equivalence of formulae modulo the class $\mathscr{K}$. To get rid of an arbitrary chain of quantifiers, we use induction on the number of quantifiers, noting that the complement of a set defined by $(\forall\zeta)\,\psi(\xi,\zeta)$ is the set defined by $(\exists\zeta)\neg\psi(\xi,\zeta)$, where $\neg(\cdot)$ denotes negation.

For some alternative proofs, see [57, Corollary A.5.2] or [67, Theorem 1.6]. □

2.3 Łoś' Theorem

Thanks to Quantifier Elimination (Theorem 2.2.1), when dealing with algebraically closed fields, we may forget altogether about formulae and use constructible subsets instead. However, we will not always be able to work just in algebraically closed fields and so we need to formulate a general transfer principle for ultraproducts. For most of our purposes, the following version suffices:

Theorem 2.3.1 (Equational Łoś' Theorem). *Suppose each A_w is an R-algebra, and let $A_\natural$ denote their ultraproduct. Let ξ be an n-tuple of variables, let $f \in R[\xi]$, and let $\mathbf{a}_w$ be n-tuples in A_w with ultraproduct $\mathbf{a}_\natural$. Then $f(\mathbf{a}_\natural) = 0$ in $A_\natural$ if and only if $f(\mathbf{a}_w) = 0$ in A_w for almost all w.*

Moreover, instead of a single equation $f = 0$, we may take in the above statement any system of equations and negations of equations over R.

Proof. Let me only sketch a proof of the first assertion. Suppose $f(\mathbf{a}_\natural) = 0$. One checks (do this!), making repeatedly use of (2.1), that $f(\mathbf{a}_\natural)$ is equal to the ultraproduct of the $f(\mathbf{a}_w)$. Hence the former being zero simply means that almost all $f(\mathbf{a}_w)$ are zero. The converse is proven by simply reversing this argument. □

On occasion, we might also want to use the full version of Łoś' Theorem, which requires the notion of a formula as defined above. Recall that a sentence is a formula without free variables.

Theorem 2.3.2 (Łoś' Theorem). *Let R be a ring and let A_w be R-algebras. If φ is a sentence with parameters from R, then φ holds in almost all A_w if and only if φ holds in the ultraproduct $A_\natural$.*

More generally, let $\varphi(\xi_1,\dots,\xi_n)$ be a formula with parameters from R and let $\mathbf{a}_w$ be an n-tuple in A_w with ultraproduct $\mathbf{a}_\natural$. Then $\varphi(\mathbf{a}_w)$ holds in almost all A_w if and only if $\varphi(\mathbf{a}_\natural)$ holds in $A_\natural$.

The proof is tedious but not hard; one simply has to unwind the definition of formula (see [57, Theorem 9.5.1] for a more general treatment). Note that $A_\natural$ is naturally an R-algebra, so that it makes sense to assert that φ is true or false in $A_\natural$. Applying Łoś' Theorem to a quantifier free formula proves Theorem 2.3.1.

2.4 Ultra-rings

An *ultra-ring* is simply an ultraproduct of rings. Probably the first examples of ultra-rings appearing in the literature are the so-called *non-standard integers*, that is to say, the ultrapowers $\mathbb{Z}_\natural$ of $\mathbb{Z}$,[2] and the *hyper-reals*, that is to say, the ultrapower $\mathbb{R}_\natural$ of the reals, which figure prominently in *non-standard analysis* (see, for instance, [36, 80]). Ultra-rings will be our main protagonists, but for the moment we only establish some very basic facts about them.

2.4.1 Ultra-fields

Let K_w be a collection of fields and $K_\natural$ their ultraproduct, which is again a field by 2.1.5 (or by an application of Łoś' Theorem). Any field which arises in this way is called an *ultra-field*.[3] Since an ultraproduct is either finite or uncountable, $\mathbb{Q}$ is an example of a field which is not an ultra-field.

2.4.1 *If for each prime number p, only finitely many K_w have characteristic p, then $K_\natural$ has characteristic zero.*

Indeed, for every prime number p, the equation $p\xi - 1 = 0$ has a solution in all but finitely many of the K_w and hence it has a solution in $K_\natural$, by Theorem 2.3.1. We will call an ultra-field $K_\natural$ of characteristic zero which arises as an ultraproduct of fields of positive characteristic, a *Lefschetz field* (the name is inspired by Theorem 2.4.3 below); and more generally, an ultra-ring of characteristic zero given as the ultraproduct of rings of positive characteristic will be called a *Lefschetz ring* (see §7.2.1 for more).

---

[2] Logicians study these under the guise of *models of Peano arithmetic*, where, instead of $\mathbb{Z}_\natural$, one traditionally looks at the sub-semi-ring $\mathbb{N}_\natural$, the ultrapower of $\mathbb{N}$ (see, for instance, [63]). Caveat: not all non-standard models are realizable as ultrapowers.

[3] In case the K_w are finite but of unbounded cardinality, their ultraproduct $K_\natural$ is also called a *pseudo-finite field*; in these notes, however, we prefer the usage of the prefix *ultra-*, and so we would call such fields instead *ultra-finite fields*.

2.4.2 *If almost all K_w are algebraically closed fields, then so is $K_\natural$.*

The quickest proof is by means of Łoś' Theorem, although one could also give an argument using just Theorem 2.3.1.

Proof. For each $n \geq 2$, consider the sentence σ_n given by

$$(\forall \zeta_0, \ldots, \zeta_n)\,(\exists \xi)\,\zeta_n = 0 \;\vee\; \zeta_n \xi^n + \cdots + \zeta_1 \xi + \zeta_0 = 0.$$

This sentence is true in any algebraically closed field, whence in almost all K_w, and therefore, by Łoś' Theorem, in $K_\natural$. However, a field in which every σ_n holds is algebraically closed. □

We have the following important corollary which can be thought of as a model-theoretic Lefschetz Principle (here $\mathbb{F}_p^{\mathrm{alg}}$ is the algebraic closure of the p-element field $\mathbb{F}_p$; and, more generally, F^{alg} denotes the algebraic closure of a field F).

Theorem 2.4.3 (Lefschetz Principle). *Let W be the set of prime numbers, endowed with some ultrafilter. The ultraproduct of the fields $\mathbb{F}_p^{alg}$ is isomorphic with the field $\mathbb{C}$ of complex numbers, that is to say, we have an isomorphism*

$$\mathbb{C} \cong \operatorname*{ulim}_{p \to \infty} \mathbb{F}_p^{alg}$$

Proof. Let $\mathbb{F}_\natural$ denote the ultraproduct of the fields $\mathbb{F}_p^{\mathrm{alg}}$. By 2.4.2, the field $\mathbb{F}_\natural$ is algebraically closed, and by 2.4.1, its characteristic is zero. Using elementary set theory, one calculates that the cardinality of $\mathbb{F}_\natural$ is equal to that of the continuum. The theorem now follows since any two algebraically closed fields of the same uncountable cardinality and the same characteristic are (non-canonically) isomorphic by Steinitz's Theorem (see [57] or Theorem 2.4.7 below). □

Remark 2.4.4. We can extend the above result as follows: any algebraically closed field K of characteristic zero and cardinality 2^κ, for some infinite cardinal κ, is a Lefschetz field. Indeed, for each p, choose an algebraically closed field K_p of characteristic p and cardinality κ. Since the ultraproduct of these fields is then an algebraically closed field of characteristic zero and cardinality 2^κ, it is isomorphic to K by Steinitz's Theorem (Theorem 2.4.7). Under the generalized Continuum Hypothesis, any uncountable cardinal is of the form 2^κ, and hence any uncountable algebraically closed field of characteristic zero is then a Lefschetz field. We will tacitly assume this, but the reader can check that nowhere this assumption is used in an essential way.

Remark 2.4.5. Theorem 2.4.3 is an embodiment of a well-known heuristic principle in algebraic geometry regarding transfer between positive and zero characteristic, which Weil [113] attributes to Lefschetz. Essentially metamathematical in nature, there have been some attempts to formulate this principle in a formal, model-theoretic language in [10,28]; for a more general version than ours, see [32, Theorem 8.3]. In fact, Theorem 2.4.3 is a special instance of model-theoretic compactness applied to the theory of algebraically closed fields. For instance, the next result, due to Ax [8], is normally proven using compactness, but here is a proof using Theorem 2.4.3 instead:

2.4.6 *If a polynomial map $\mathbb{C}^n \to \mathbb{C}^n$ is injective, then it is surjective.*

Indeed, by the Pigeon Hole Principle, the result is true if we replace $\mathbb{C}$ by any finite field; since $\mathbb{F}_p^{\mathrm{alg}}$ is a union of finite fields, the assertion remains true over it; an application of Theorem 2.4.3 then finishes the proof. □

Theorem 2.4.7 (Steinitz's Theorem). *If K and L are algebraically closed fields of the same characteristic and the same uncountable cardinality, then they are isomorphic.*

Proof (Sketch of proof). Let k be the common prime field of K and L (that is to say, either $\mathbb{Q}$ in characteristic zero, or $\mathbb{F}_p$ in positive characteristic p). Let Γ and Δ be respective transcendence bases of K and L over k. Since K and L have the same uncountable cardinality, Γ and Δ have the same cardinality, and hence there exists a bijection $f\colon \Gamma \to \Delta$. This naturally extends to a field isomorphism $k(\Gamma) \to k(\Delta)$. Since K is the algebraic closure of $k(\Gamma)$, and similarly, L of $k(\Delta)$, this isomorphism then extends to an isomorphism $K \to L$. □

The previous results might lead the reader to think that the choice of ultrafilter never matters. As we shall see later, for most of our purposes this is indeed true, but there are many situations were the ultrafilter determines the ultraproduct. For instance, consider the ultraproduct of fields F_w, where F_w is either $\mathbb{F}_2$ or $\mathbb{F}_3$. Since almost all F_w are therefore equal to one, and only one, among these two fields, so will their ultraproduct be (to see the latter, note that there is a first-order sentence expressing that a field has exactly two elements, and now use the model-theoretic version of Łoś' Theorem, Theorem 2.3.2). More precisely, the ultraproduct is equal to $\mathbb{F}_2$ if and only if the set I_2 of indices w for which $F_w = \mathbb{F}_2$ belongs to the ultrafilter. If I_2 is infinite, then there exists always an ultrafilter containing it, and if I_2 is also co-finite, then there exists another one not containing I_2, so that in the former case, the ultraproduct is equal to $\mathbb{F}_2$, and in the latter case to $\mathbb{F}_3$. We will prove a theorem below, Theorem 2.5.4, which tells us exactly all possible ultraproducts a given collection of rings can produce (see also Theorem 2.6.4).

2.4.2 Ultra-rings

Let A_w be a collection of rings. Their ultraproduct $A_\natural$ will be called, as already mentioned, an *ultra-ring*.

2.4.8 *If each A_w is local with maximal ideal $\mathfrak{m}_w$ and residue field $k_w := A_w/\mathfrak{m}_w$, then $A_\natural$ is local with maximal ideal $\mathfrak{m}_\natural := \operatorname{ulim} \mathfrak{m}_w$ and residue field $k_\natural := \operatorname{ulim} k_w$.*

Indeed, a ring is local if and only if the sum of any two non-units is again a non-unit. This statement is clearly expressible by means of a sentence, so that by Łoś' Theorem (Theorem 2.3.2), $A_\natural$ is local. Again we can prove this also directly, or using the equational version, Theorem 2.3.1. The remaining assertions now follow easily from 2.1.6. In fact, the same argument shows that the converse is also true: if $A_\natural$ is local, then so are almost all A_w. □

Recall that the *embedding dimension* of a local ring is the minimal number of generators of its maximal ideal. The next result is therefore immediate from 2.1.6 and 2.4.8.

2.4.9 *If A_w are local rings of embedding dimension e, then so is $A_\natural$.* □

As being a domain is captured by the fact that the equation $\xi\zeta = 0$ has no solution by non-zero elements; and being reduced by the fact that the equation $\xi^2 = 0$ has no non-zero solutions, we immediately get from Łoś' Theorem:

2.4.10 *Almost all A_w are domains (respectively, reduced) if and only if $A_\natural$ is a domain (respectively, reduced).* □

In particular, using 2.1.6, we see that an ultraproduct of ideals is a prime (respectively, radical, maximal) ideal if and only if almost all ideals are prime (respectively, reduced, maximal).

2.4.11 *If I_w are ideals in the local rings $(A_w, \mathfrak{m}_w)$, such that in $(A_\natural, \mathfrak{m}_\natural)$, their ultraproduct $I_\natural$ is $\mathfrak{m}_\natural$-primary, then almost all I_w are $\mathfrak{m}_w$-primary.*

Recall that an ideal I in a local ring $(R, \mathfrak{m})$ is called $\mathfrak{m}$-*primary* if its radical is equal to $\mathfrak{m}$. So, $\mathfrak{m}_\natural^N \subseteq I_\natural$ for some N, and therefore, $\mathfrak{m}_w \subseteq I_w$ for almost all w, by Łoś' Theorem. □

Note that here the converse may fail to hold: not every ultraproduct of $\mathfrak{m}_w$-primary ideals need to be $\mathfrak{m}_\natural$-primary (see Proposition 2.4.17 for a partial converse). For instance, the ultraproduct of the $\mathfrak{m}^w$ is no longer $\mathfrak{m}R_\natural$-primary in the ultrapower $R_\natural$ (see 8.1.3). An ideal in an ultra-ring is called an *ultra-ideal*, if it is an ultraproduct of ideals.[4]

2.4.12 *Any finitely generated, or more generally, any finitely related ideal $\mathfrak{a}$ in an ultra-ring $A_\natural$ is an ultra-ideal, and $A_\natural/\mathfrak{a}$ is again an ultra-ring.*

Let $A_\natural$ be the ultraproduct of rings A_w. Recall that an ideal $\mathfrak{a}$ is called *finitely related*, if it is of the form $(I : J)$ with I and J finitely generated. Suppose $I = (f_1, \dots, f_n)A_\natural$ and $J = (g_1, \dots, g_m)A_\natural$. Choose $f_{iw}, g_{iw} \in A_w$ with ultraproduct equal to f_i and g_i respectively, and put

$$\mathfrak{a}_w := ((f_{1w}, \dots, f_{nw})A_w : (g_{1w}, \dots, g_{mw})A_w).$$

[4] In the literature, such ideals are often called *internal* ideals.

It is now an easy exercise on Łoś' Theorem, using 2.1.6, that $\mathfrak{a}$ is the ultraproduct of the $\mathfrak{a}_w$, and $A_\natural/\mathfrak{a}$ the ultraproduct of the $A_w/\mathfrak{a}_w$. □

Not every ideal in an ultra-ring is an ultra-ideal; for an example, see the discussion at the start of §4.2. Another counterexample is provided by the following ideal, which will play an important role in the study of local ultra-rings (see Proposition 2.4.19 for an example).

Definition 2.4.13 (Ideal of Infinitesimals). For an arbitrary local ring $(R,\mathfrak{m})$, define its *ideal of infinitesimals*, denoted $\mathfrak{I}_R$, as the intersection

$$\mathfrak{I}_R := \bigcap_{n\geq 0} \mathfrak{m}^n.$$

The $\mathfrak{m}$-adic topology on R is Hausdorff (=separated) if and only if $\mathfrak{I}_R = 0$. Therefore, we will refer to the residue ring $R/\mathfrak{I}_R$ as the *separated quotient* of R. In commutative algebra, the ideal of infinitesimals hardly ever appears simply because of:

Theorem 2.4.14 (Krull's Intersection Theorem). *If R is a Noetherian local ring, then $\mathfrak{I}_R = 0$.*

Proof. This is an immediate consequence of the Artin-Rees Lemma (for which see [69, Theorem 8.5] or [7, Proposition 10.9]), or of its weaker variant proven in Theorem 8.2.1 below. Namely, for $x \in \mathfrak{I}_R$, there exists, according to the latter theorem, some c such that $xR \cap \mathfrak{m}^c \subseteq x\mathfrak{m}$. Since $x \in \mathfrak{m}^c$ by assumption, we get $x \in x\mathfrak{m}$, that is to say, $x = ax$ with $a \in \mathfrak{m}$. Hence $(1-a)x = 0$. As $1-a$ is a unit in R, we get $x = 0$. □

It would be dishonest to claim that the above yields a non-standard proof of Krull's theorem via Theorem 8.2.1, as the latter proof uses the flatness of cataproducts (Theorem 8.1.15), which is obtained via Cohen's Structure Theorems, and therefore, ultimately relies on Krull's Intersection Theorem. The exact connection between both results is given by Theorem 8.2.3.

Corollary 2.4.15. *In a Noetherian local ring $(R,\mathfrak{m})$, every ideal is the intersection of $\mathfrak{m}$-primary ideals.*

Proof. For $I \subseteq R$ an ideal, an application of Theorem 2.4.14 to the ring R/I shows that I is the intersection of all $I + \mathfrak{m}^n$, and the latter are indeed $\mathfrak{m}$-primary. □

Most local ultra-rings have a non-zero ideal of infinitesimals.

2.4.16 *If R_w are local rings with non-nilpotent maximal ideal, then the ideal of infinitesimals of their ultraproduct $R_\natural$ is non-zero. In particular, $R_\natural$ is not Noetherian.*

Indeed, by assumption, we can find non-zero $a_w \in \mathfrak{m}^w$ (let us for the moment assume that the index set is equal to $\mathbb{N}$) for all w. Hence their ultraproduct $a_\natural$ is non-zero and lies inside $\mathfrak{I}_{R_\natural}$. □

As we shall see later, being Noetherian is not preserved under ultraproducts. However, under certain restrictive conditions, of which the field case (2.1.5) is a special instance, we do have preservation (this also gives a more quantitative version of 2.4.11):

Proposition 2.4.17. *An ultraproduct $A_\natural$ of rings A_w is Artinian of length l if and only if almost all A_w are Artinian of length l.*

Proof. By the Jordan-Holder Theorem, there exist elements $a_0 = 0, a_1, \dots, a_l = 1$ in $A_\natural$ such that

$$a_0 A_\natural \subsetneq (a_0, a_1)A_\natural \subsetneq (a_0, a_1, a_2)A_\natural \subsetneq \cdots \subsetneq (a_0, \dots, a_l)A_\natural = A_\natural$$

is a maximal chain of ideals. Choose, for each $i = 0, \dots, l$, elements $a_{iw} \in A_w$ whose ultraproduct is a_i. By Łoś' Theorem, for a fixed $i < l$, almost all inclusions

$$(a_{0w}, \dots, a_{iw})A_w \subseteq (a_{0w}, \dots, a_{i+1w})A_w \tag{2.3}$$

are strict. This shows that almost all A_w have length at least l. If almost all of them would have length bigger than l, then for at least one i, we can insert in almost all inclusions (2.3) an ideal I_w different from both ideals. By Łoś' Theorem, the ultraproduct $I_\natural$ of the I_w would then be strictly contained between $(a_0, \dots, a_i)A_\natural$ and $(a_0, \dots, a_{i+1})A_\natural$, implying that $A_\natural$ has length at least $l + 1$, contradiction. □

Proposition 2.4.18. *An* ultra-Dedekind domain*, that is to say, an ultraproduct of Dedekind domains, is a Prüfer domain.*

Proof. Recall that a domain is *Prüfer* if any localization at a maximal ideal is a valuation ring. By [34, §1.4], this equivalent with the property that every finitely generated ideal is projective, and so we verify the latter. Let A_w be Dedekind domains, that is to say, one-dimensional normal domains, and let $A_\natural$ be their ultraproduct. Let $I_\natural$ be a finitely generated ideal. By 2.4.12, we can find non-zero ideals $I_w \subseteq A_w$ such that their ultraproduct equals $I_\natural$. Since each I_w is generated by at most two elements we can find a split exact sequence

$$0 \to J_w \to A_w^2 \to I_w \to 0$$

for some submodule $J_w \subseteq A_w^2$. Since ultraproducts commute with direct sums, we get an isomorphism $I_\natural \oplus J_\natural \cong A_\natural^2$, where $J_\natural$ is the ultraproduct of the J_w, showing that $I_\natural$ is projective. □

Proposition 2.4.19. *An* ultra-discrete valuation ring $V_\natural$*, that is to say, an ultraproduct of discrete valuation rings V_w, is a valuation domain. Its ideal of infinitesimals $\mathfrak{I}_{V_\natural}$ is an infinitely generated prime ideal, and the separated quotient $V_\natural/\mathfrak{I}_{V_\natural}$—in Chapter 8 we will call this the* cataproduct $V_\sharp$ *of the V_w—is again a discrete valuation ring.*

Proof. Recall that a *valuation ring* is a domain such that for all a in the field of fractions of V, at least one of a or $1/a$ belongs to V. By 2.1.5, the field of fractions $K_\natural$ of $V_\natural$ is the ultraproduct of the field of fractions K_w of the V_w. Let $a_w \in K_w$ be an approximation of $a \in K_\natural$. For almost each w, either a_w or $1/a_w$ belongs to V_w. Therefore, by Łoś' Theorem, either $a \in V_\natural$ or $1/a \in V_\natural$, proving the first claim. If $\mathfrak{I}_{V_\natural}$ is finitely generated, then it is principal, say, of the form $b_\natural V_\natural$, since $V_\natural$ is a valuation domain. Let $b_w \in V_w$ be an approximation of $b_\natural$, and let $c_w := b_w/\pi_w$, where π_w is a uniformizing parameter of V_w. Since, for each n, almost all b_w have order at least n, almost all c_w have order at least $n-1$. Hence their ultraproduct $c_\natural$ also belongs to $\mathfrak{I}_{V_\natural} = b_\natural V_\natural$. Let $\pi_\natural$ be the ultraproduct of the π_w, so that it generates the maximal ideal $\mathfrak{m}_\natural$ of $V_\natural$ by the proof of 2.4.9. By Łoś' Theorem, $b_\natural/\pi_\natural = c_\natural$, so that $b_\natural \in c_\natural \mathfrak{m}_\natural \subseteq b_\natural \mathfrak{m}_\natural$, contradiction. Finally, to show that $V_\sharp := V_\natural/\mathfrak{I}_{V_\natural}$ is a discrete valuation ring, and hence, in particular, $\mathfrak{I}_{V_\natural}$ is prime, observe that for any non-zero element a in $V_\sharp$, there is a largest n such that $a \in \mathfrak{m}_\natural^n = \pi_\natural^n V_\natural$. The assignment $a \mapsto n$ is now easily seen to be a discrete valuation. □

The previous proof in fact shows that an ultraproduct of valuation rings is again a valuation ring.

2.4.3 Ultrapowers

An important instance of an ultra-ring is the ultrapower $A_\natural$ of a ring A. It is easy to see that the diagonal embedding $A \to A_\natural$ is a ring homomorphism. We will see in the next chapter that this embedding is often flat (see Corollary 3.3.3 and Theorem 3.3.4). However, an easy application of Łoś' Theorem immediately yields that this map is at least cyclically pure. Recall that a homomorphism $A \to B$ is called *cyclically pure*, if $IB \cap A = I$ for all ideals $I \subseteq A$. Examples of cyclically pure homomorphisms are, as we shall see, faithfully flat (Proposition 3.2.5) and split maps (see 5.5.4). It follows from Proposition 2.4.17 that the ultrapower of an Artinian ring is again Artinian. However, by 2.4.16 and Theorem 2.4.14, these are the only rings whose ultrapower is Noetherian. The next result is immediate from 2.1.6 and its proof:

2.4.20 *If I is a finitely generated ideal in a ring A, then its ultrapower in the ultrapower $A_\natural$ of A is equal to $IA_\natural$. In particular, the ultrapower of A/I is $A_\natural/IA_\natural$.*

The following is a counterexample if I is not finitely generated: let A be the polynomial ring over a field in countably many variables ξ_i, and let I be the ideal generated by all these variables. The ultraproduct f of the polynomials $f_w := \xi_1 + \cdots + \xi_w$ is an element in the ultrapower $I_\natural$ of I but does not belong to $IA_\natural$, for if it were, then f must be a sum of finitely many generators of I, say, $\xi_1, \dots, \xi_i$, and therefore by 2.1.6, so must almost all f_w be, a contradiction whenever $w > i$.

2.4.4 Ultra-exponentiation

Let $A_\natural$ be an ultra-ring, given as the ultraproduct of rings A_w. Let $\mathbb{N}_\natural$ be the ultrapower of the natural numbers, and let $\alpha \in \mathbb{N}_\natural$ with approximations α_w. The *ultra-exponentiation map* on A with exponent α is defined as follows. Given $x \in A$, let $x_w \in A_w$ be an approximation of x, that is to say, $\operatorname{ulim} x_w = x$, and set

$$x^\alpha := \operatorname{ulim} x_w^{\alpha_w}.$$

One easily verifies that this definition does not depend on the choice of approximation of x nor of α: if x'_w and α'_w are also respectively approximations of x and α, then almost all x_w and x'_w are the same, and so are almost all α_w and α'_w, whence almost all $x_w^{\alpha_w}$ are equal to $(x'_w)^{\alpha'_w}$, and, therefore, they have the same ultraproduct. By Łoś' Theorem, ultra-exponentiation satisfies the same rules as regular exponentiation:

$$(xy)^\alpha = x^\alpha \cdot y^\alpha \quad \text{and} \quad x^\alpha \cdot x^\beta = x^{\alpha+\beta} \quad \text{and} \quad (x^\alpha)^\beta = x^{\alpha\cdot\beta}$$

for all $x, y \in A_\natural$ and all $\alpha, \beta \in \mathbb{N}_\natural$.

If A is local and x a non-unit, then x^α is an infinitesimal for any α in $\mathbb{N}_\natural$ not in $\mathbb{N}$. In these notes, the most important instance will be the ultra-exponentiation map obtained as the ultraproduct of Frobenius maps. More precisely, let $A_\natural$ be a Lefschetz ring, say, realized as the ultraproduct of rings A_p of characteristic p (here we assumed for simplicity that the underlying index set is just the set of prime numbers, but this is not necessary). On each A_p, we have an action of the *Frobenius*, given as $\mathbf{F}_p(x) := x^p$ (for more, see §5.1).

Definition 2.4.21 (Ultra-Frobenius). The ultraproduct of these Frobenii yields an endomorphism $\mathbf{F}_\natural$ on $A_\natural$, called the *ultra-Frobenius*, given by $\mathbf{F}_\natural(x) := x^\pi$, where $\pi \in \mathbb{N}_\natural$ is the ultraproduct of all prime numbers. Since each Frobenius is an endomorphism, so is any ultra-Frobenius by 2.1.7. In particular, we have

$$(x+y)^\pi = x^\pi + y^\pi$$

for all $x, y \in A_\natural$.

2.5 Algebraic Definition of Ultra-rings

Let A_w, for $w \in W$, be rings with Cartesian product $A_\infty := \prod_w A_w$ and direct sum $A_{(\infty)} := \bigoplus A_w$. Note that $A_{(\infty)}$ is an ideal in A_∞. Call an element $a \in A_\infty$ a *strong idempotent* if each of its entries is either zero or one. In other words, an element in A_∞ is a strong idempotent if and only if it is the characteristic function 1_D of a subset $D \subseteq W$. For any ideal $\mathfrak{a} \subseteq A_\infty$, let $\mathfrak{a}^\circ$ be the ideal generated by all strong idempotents in $\mathfrak{a}$, and let $\mathfrak{W}_\mathfrak{a}$ be the collection of subsets $D \subseteq W$ such that $1-1_D \in \mathfrak{a}$. Using the identities $(1-1_D)(1-1_E) = 1-1_E$ for $D \subseteq E$ and $1-1_{D\cap E} = 1_E(1-1_D)+1-1_E$, one verifies that $\mathfrak{W}_\mathfrak{a}$ is a filter.

2.5.1 *Given an ideal $\mathfrak{a} \subseteq A_\infty$, the filter $\mathfrak{W}_\mathfrak{a}$ is maximal if and only if $\mathfrak{a}$ is a prime ideal; it is principal if and only if the ideal $\mathfrak{a}^\circ$ is principal, if and only if $\mathfrak{a}$ does not contain the ideal $A_{(\infty)}$.*

Indeed, given an idempotent e, its *complement* $1-e$ is again idempotent, and the product of both is zero, that is to say, they are orthogonal. It follows that any prime ideal contains exactly one among e and $1-e$. Hence, if $\mathfrak{a}$ is prime, then $\mathfrak{W}_\mathfrak{a}$ consists of those subsets $D \subseteq W$ such that $1_D \notin \mathfrak{a}$. Since $1-1_D$ is the characteristic function of the complement of D, it follows that either D or its complement belongs to $\mathfrak{W}_\mathfrak{a}$. Moreover, if $D \in \mathfrak{W}_\mathfrak{a}$ and $D \subseteq E$, then $1_D \cdot 1_E = 1_D$ does not belong to $\mathfrak{a}$, whence neither does 1_E, showing that $E \in \mathfrak{W}_\mathfrak{a}$. This proves that $\mathfrak{W}_\mathfrak{a}$ is a maximal filter. It is not hard to see that if $\mathfrak{a}^\circ$ is principal, then it must be generated by the characteristic function of the complement of a singleton, and hence $\mathfrak{W}_\mathfrak{a}$ must be principal (the other direction is immediate). The last equivalence is left as an exercise to the reader. □

We can now formulate the following entirely algebraic characterization of an ultra-ring.

2.5.2 *Let $\mathfrak{P}$ be a prime ideal of A_∞ containing the direct sum ideal $A_{(\infty)}$. The ultraproduct of the A_w with respect to the ultrafilter $\mathfrak{W}_\mathfrak{P}$ is equal to $A_\infty/\mathfrak{P}^\circ$, that is to say, $\mathfrak{P}^\circ$ is the null-ideal determined by $\mathfrak{W}_\mathfrak{P}$. Furthermore, any ultra-ring having the A_w as approximations is of the form $A_\infty/\mathfrak{P}^\circ$, for some prime ideal $\mathfrak{P}$ containing $A_{(\infty)}$.*

Let $\mathfrak{n}$ be the null-ideal determined by $\mathfrak{W}_\mathfrak{P}$, that is to say, the collection of sequences in A_∞ almost all of whose entries are zero. If $D \in \mathfrak{W}_\mathfrak{P}$, then almost all entries of $1-1_D$ are zero, and hence $1-1_D \in \mathfrak{n}$. Since this is a typical generator of $\mathfrak{P}^\circ$, we get $\mathfrak{P}^\circ \subseteq \mathfrak{n}$. Conversely, suppose $a = (a_w) \in \mathfrak{n}$. Hence $a_w = 0$ for all w belonging to some $D \in \mathfrak{W}_\mathfrak{P}$. Since $1-1_D \in \mathfrak{P}^\circ$ and $a = a(1-1_D)$, we get $a \in \mathfrak{P}^\circ$.

Conversely, if $\mathfrak{W}$ is an ultrafilter with corresponding null-ideal $\mathfrak{n} \subseteq A_\infty$, then one easily checks that any prime ideal $\mathfrak{P}$ containing $\mathfrak{n}$ satisfies $\mathfrak{n} = \mathfrak{P}^\circ$. □

In fact, if $\mathfrak{P} \subseteq \mathfrak{Q}$ are prime ideals, then $\mathfrak{P}^\circ = \mathfrak{Q}^\circ$, showing that already all minimal prime ideals of A_∞ determine all possible ultrafilters.

Corollary 2.5.3. *If all A_w are domains, then $A_\natural$ is the coordinate ring of an irreducible component of* $\operatorname{Spec}(A_\infty/A_{(\infty)})$. *More precisely, the residue rings $A_\infty/\mathfrak{G}$, for $\mathfrak{G} \subseteq A_\infty$ a minimal prime containing $A_{(\infty)}$, are precisely the ultraproducts $A_\natural$ having the domains A_w for approximations. Moreover, these irreducible components are then also the connected components of* $\operatorname{Spec}(A_\infty/A_{(\infty)})$, *that is to say, they are mutually disjoint.*

Proof. Since the ultraproduct $A_\natural$ determined by $\mathfrak{G}$ is equal to $A_\infty/\mathfrak{G}^\circ$ by 2.5.2, and a domain by 2.4.10, the ideal $\mathfrak{G}^\circ$ must be prime. By minimality, $\mathfrak{G}^\circ = \mathfrak{G}$.

To prove the last assertion, let $\mathfrak{G}_1$ and $\mathfrak{G}_2$ be two distinct minimal prime ideals of A_∞ containing $A_{(\infty)}$. Suppose $\mathfrak{G}_1 + \mathfrak{G}_2$ is not the unit ideal. Hence there exists a maximal ideal $\mathfrak{M} \subseteq A_\infty$ such that $\mathfrak{G}_1, \mathfrak{G}_2 \subseteq \mathfrak{M}$, and hence

$$\mathfrak{G}_1 = \mathfrak{G}_1^\circ = \mathfrak{M}^\circ = \mathfrak{G}_2^\circ = \mathfrak{G}_2,$$

contradiction. Hence $\mathfrak{G}_1 + \mathfrak{G}_2 = 1$, showing that any two irreducible components of $\operatorname{Spec}(A_\infty/A_{(\infty)})$ are disjoint. □

Note that the connected components of $\operatorname{Spec}(A_\infty)$, apart from the $\operatorname{Spec}(A_\natural)$, are the $\operatorname{Spec}(A_w)$ corresponding to the principal (maximal) filters. In the following structure theorem, $\mathbb{Z}_\infty := \mathbb{Z}^W$ denotes the Cartesian power of $\mathbb{Z}$. Any Cartesian product $A_\infty := \prod A_w$ is naturally a $\mathbb{Z}_\infty$-algebra.

Theorem 2.5.4. *Any ultra-ring is a base change of a ring of non-standard integers $\mathbb{Z}_\natural$. More precisely, the ultra-rings with approximation A_w are precisely the rings of the form $A_\infty/\mathfrak{G}A_\infty$, where $\mathfrak{G}$ is a minimal prime of $\mathbb{Z}_\infty$ containing the direct sum ideal.*

Proof. If $\mathfrak{P}$ is a prime ideal in A_∞ containing the direct sum ideal $A_{(\infty)}$, then the generators of $\mathfrak{P}^\circ$ already live in $\mathbb{Z}_\infty$, and generate the null-ideal in $\mathbb{Z}_\infty$ corresponding to the ultrafilter $\mathfrak{W}_\mathfrak{P}$. By Corollary 2.5.3, the latter ideal therefore is a minimal prime ideal $\mathfrak{G} \subseteq \mathbb{Z}_\infty$ of $\mathbb{Z}_{(\infty)}$. Since $\mathfrak{G}A_\infty = \mathfrak{P}^\circ$, one direction is clear from 2.5.2. Conversely, again by Corollary 2.5.3, any minimal prime ideal $\mathfrak{G} \subseteq \mathbb{Z}_\infty$ is the null-ideal determined by the ultrafilter $\mathfrak{W}_\mathfrak{G}$, and one easily checks that the same is therefore true for its extension $\mathfrak{G}A_\infty$. □

2.6 Sheaf-theoretic Definition of Ultra-rings

We say that a topological Hausdorff space X admits a *Hausdorff compactification* $X^\vee$, if $X \subseteq X^\vee$ such that for every compact Hausdorff space Y and every continuous map $f\colon X \to Y$, there is a unique map $f^\vee\colon X^\vee \to Y$ extending f. Since this is a universal problem, a Hausdorff compactification is unique, if it exists.

Proposition 2.6.1. *Every infinite discrete space X has a Hausdorff compactification.*

Proof. Let $X^\vee$ be the *Stone-Čech compactification* of X consisting of all maximal filters on X. We identify the principal filters with their generators, so that X becomes a subset of $X^\vee$. For a subset $U \subseteq X$, let $\tau(U) \subseteq X^\vee$ consist of all maximal filters containing U. For any $U \subseteq X$, we have

$$X^\vee - \tau(U) = \tau(X - U), \tag{2.4}$$

by the ultrafilter condition. We define a topology on X by taking the $\tau(U)$, for $U \subseteq X$, as a basis of open subsets. This works, since the intersection of two basic opens $\tau(U_1)$ and $\tau(U_2)$ is the basic open $\tau(U_1 \cap U_2)$. Note that $U \subseteq \tau(U)$ with equality if and only if U is finite. In fact, by the definition of the embedding $X \subseteq X^\vee$, we have $\tau(U) \cap X = U$, and hence the topology induced on X is just the discrete topology. In particular, every non-empty open has a non-empty intersection with X, showing that X is a dense (open) subset of $X^\vee$.

To see that $X^\vee$ is Hausdorff, take two distinct points in $X^\vee$, that is to say, distinct maximal filters on X. In particular, there exists a subset $U \subseteq X$ belonging to one but not the other. Hence $\tau(U)$ and $\tau(X - U)$ are disjoint opens, each containing exactly one of these two points. To prove compactness, we need to verify that the finite intersection property holds, that is to say, that any collection of non-empty closed subsets which is closed under finite intersections, has non-empty intersection. By (2.4), any basic open subset is also closed, that is to say, is a *clopen*, and hence any closed subset is an intersection of basic opens. Without loss of generality, we may therefore assume that $\{\tau(U_i)\}_i$ is a collection of non-empty closed subsets which is closed under finite intersections, and we have to show that their intersection is also non-empty. Since $X \cap \tau(U_i) = U_i$, it follows that the $\{U_i\}_i$ are closed under finite intersections. Let $\mathfrak{Y}$ be the collection of all subsets $U \subseteq X$ such that some U_i is contained in U. One checks that $\mathfrak{Y}$ is a filter, whence is contained in some maximal filter $\mathfrak{W}$. By construction, $U_i \in \mathfrak{W}$, for all i, showing that $\mathfrak{W}$ lies in the intersection of all $\tau(U_i)$.

Finally, we verify the universal property. Let $f\colon X \to Y$ be an (automatically continuous) map with Y a compact Hausdorff space and fix a point in $X^\vee$, that is to say, a maximal filter $\mathfrak{W}$ on X. Let $F_{\mathfrak{W}}$ be the intersection of all closures $\operatorname{clos}(f(U))$, where U runs over all subsets in $\mathfrak{W}$. Since any finite intersection

$$\operatorname{clos}(f(U_1)) \cap \cdots \cap \operatorname{clos}(f(U_s)),$$

for $U_i \in \mathfrak{W}$ contains the (non-empty) image of $U_1 \cap \cdots \cap U_s \in \mathfrak{W}$ under f, and since Y is compact, $F_{\mathfrak{W}}$ is non-empty. Suppose y and y' are two distinct elements in $F_{\mathfrak{W}}$. Since Y is Hausdorff, we can find disjoint opens V and V' containing respectively y and y'. In particular, their pre-images $f^{-1}V$ and $f^{-1}V'$ are disjoint, and so one of them, say $f^{-1}V$ cannot belong to $\mathfrak{W}$. It follows that $X - f^{-1}V$ belongs to $\mathfrak{W}$ and hence $F_{\mathfrak{W}}$ is contained in the closure of $f(X - f^{-1}V) = f(X) - V$. Since V is an open containing $y \in F_{\mathfrak{W}}$, it must therefore have non-empty intersection with $f(X) - V$, contradiction. Hence $F_{\mathfrak{W}}$ is a singleton, and we now define $f^\vee(\mathfrak{W})$ to be the unique element belonging to $F_{\mathfrak{W}}$. Immediate from the definitions we get that $f^\vee(\mathfrak{W}) = f(\mathfrak{W})$ in case $\mathfrak{W} \in X$, that is to say, is principal. So remains to show that $f^\vee$ is continuous.

To this end, let $V \subseteq Y$ be open and $\mathfrak{W} \in X^\vee$ a point in $(f^\vee)^{-1}(V)$. We need to find an open containing $\mathfrak{W}$ and contained in $(f^\vee)^{-1}(V)$. By construction, the intersection of all $\operatorname{clos}(f(U))$ with $U \in \mathfrak{W}$ is contained in V. By compactness, already finitely many of the $\operatorname{clos}(f(U))$ have an intersection contained in V (since their complements together with V form an open cover of Y). Letting $U \in \mathfrak{W}$ be the intersection of these finitely many members of $\mathfrak{W}$, then, as above, $\operatorname{clos}(f(U)) \subseteq V$. To see that $\tau(U) \subseteq (f^\vee)^{-1}(V)$, take $\mathfrak{Y} \in \tau(U)$. So $U \in \mathfrak{Y}$ and hence, per construction, $f^\vee(\mathfrak{Y}) \in \operatorname{clos}(f(U)) \subseteq V$. □

If U is infinite, then intersecting each set in $\mathfrak{W} \in \tau(U)$ with U yields a maximal filter on U, so that we get an induced map $\tau(U) \to U^\vee$. It is not hard to show that this is in fact an homeomorphism. Since U is dense in $U^\vee$, we showed that the closure of U in $X^\vee$ is just $\tau(U) = U^\vee$. Let A_w be rings, indexed by $w \in X$. Define a sheaf of rings $\mathscr{A}$ on X by taking for stalk $\mathscr{A}_w := A_w$ in each point $w \in X$ (note that since X is discrete, this completely determines the sheaf $\mathscr{A}$). Let $i\colon X \to X^\vee$ be the above embedding and let $\mathscr{A}^\vee := i_*\mathscr{A}$ be the direct image sheaf of $\mathscr{A}$ under i. By general sheaf theory, this is a sheaf on $X^\vee$. For instance, on a basic open $\tau(U)$ the ring of sections of $\mathscr{A}^\vee$ is $\mathscr{A}(\tau(U) \cap X) = \mathscr{A}(U)$, and the latter is just the Cartesian product of all A_w for $w \in U$.

2.6.2 *The stalk of $\mathscr{A}^\vee$ in a boundary point $\mathfrak{W} \in X^\vee - X$ is isomorphic to the ultraproduct* $\operatorname{ulim} A_w$ *with respect to the ultrafilter $\mathfrak{W}$.*

Indeed, by definition of stalk, $\mathscr{A}^\vee_{\mathfrak{W}}$ is the direct limit of all $\mathscr{A}^\vee(V)$ where V runs over all open subsets of $X^\vee$ containing $\mathfrak{W}$. It suffices to take the direct limit over all basic opens $\tau(U)$ containing $\mathfrak{W}$, that is to say, for $U \in \mathfrak{W}$. Now, as we already observed above,

$$\mathscr{A}^\vee(\tau(U)) = \mathscr{A}(U) = \prod_{w \in U} A_w \cong A_\infty/(1 - 1_U)A_\infty.$$

Hence this direct limit is equal to the residue ring of A_∞ modulo the ideal generated by all $1 - 1_U$ for $U \in \mathfrak{W}$, that is to say, by 2.5.2, modulo the null-ideal corresponding to $\mathfrak{W}$. □

2.6.3 *Under the identification of $\mathscr{P}(X)$ with the Cartesian power $(\mathbb{F}_2)_\infty$ (see Example B.2.3), the assignment $\mathfrak{p} \mapsto \mathfrak{W}_\mathfrak{p}$ defined in §2.5 yields a homeomorphism between the affine scheme* $\operatorname{Spec}(\mathscr{P}(X))$ *and $X^\vee$. Infinitely generated prime ideals then correspond to ultrafilters.*

The only thing to observe is that the inverse image of the basic open $\tau(U)$ for $U \subseteq X$ is the basic open $D(1_{-U})$ in $\operatorname{Spec}(\mathscr{P}(X))$. Note that if we view $X^\vee$ as the set of maximal filters, and hence as a subset of $\mathscr{P}(X)$, then $\mathfrak{p}$ is sent to its complement under this homeomorphism. □

Let us call a scheme X *Boolean* if it admits an open covering by affine schemes of the form $\operatorname{Spec} B$ with B a Boolean ring (see Proposition B.1.5 for some basic properties of Boolean rings). Equivalently, any section ring is Boolean, and this is also equivalent by (B.1.5.vii) to all stalks having two elements. In particular, $\operatorname{Spec}(\mathscr{P}(W))$ is Boolean, for any set W. We call $x \in X$ a *finite point* if the prime ideal associated to x is finitely generated, whence principal by (B.1.5.iii); in the remaining case, we call x an *infinite point*. By (B.1.5.x), the infinite points form a closed subset with ideal of definition the ideal generated by all atoms. We call X *atomless* if every point is infinite, and by (B.1.5.x) this is equivalent with any section ring of an open subset being atomless. The dichotomy between finite and infinite points is robust by Corollary B.1.8, in the sense given in 2.6.5 below.

Theorem 2.6.4. *Let $\mathscr{A}$ be a sheaf of rings on a Boolean scheme X. If $x \in X$ is infinite, then the stalk $\mathscr{A}_x$ is an ultra-ring. If X is the affine scheme of a power set ring $\mathscr{P}(W)$, then $\mathscr{A}_x$ is the ultraproduct of the stalks $\mathscr{A}_y$ at finite points $y \in X$ with respect to the ultrafilter given as the image of x under the homeomorphism $X \cong W^\vee$ from 2.6.3.*

Proof. Let us first show this in the case X is $\mathrm{Spec}(\mathscr{P}(W))$. Infinite points correspond to ultrafilters by 2.6.3, and the result follows by 2.6.2. In the general case, since stalks are local objects, we may assume that X is an affine scheme with Boolean coordinate ring B. By the Stone Representation Theorem (see Theorem B.2.7 below for a proof), there exists a faithfully flat embedding $B \subseteq C := \mathscr{P}(W)$ for some W (we will actually show that one can take W equal to an ultrapower of $\mathbb{N}$). Since x is infinite, B must be infinite by (B.1.5.viii), whence so must W be. Let $Y := \mathrm{Spec}(C)$. We need:

2.6.5 *If $f\colon Y \to X$ is a dominant morphism of Boolean schemes and $x \in X$ is infinite, then there exists an infinite $y \in Y$ with $f(y) = x$.*

Indeed, we may reduce to the affine case, in which case we have an injective homomorphism $B \to C$ between Boolean rings and a non-principal maximal ideal $\mathfrak{p} \subseteq B$ corresponding to x. The fiber $f^{-1}(x)$ has coordinate ring $C/\mathfrak{p}C$. If $C/\mathfrak{p}C$ is infinite, then it contains a non-principal maximal ideal by (B.1.5.viii), and its pre-image in C must then also be non-principal, so that we are done in this case. So assume $C/\mathfrak{p}C$ is finite and any maximal ideal containing $\mathfrak{p}C$ is principal. Since $C/\mathfrak{p}C$ is finite, $\mathfrak{p}C$ is the intersection of the finitely many maximal ideals containing $\mathfrak{p}$ by (B.1.5.vi). Hence $\mathfrak{p}C$ is an intersection of principal ideals whence is principal by (B.1.5.iii). Since $B \to C$ is an embedding, $\mathfrak{p}$ must be principal by Corollary B.1.8, contradiction. □

So, returning to the case at hand, there exists an infinite $y \in Y$ such that $f(y) = x$. Let $f^{-1}\mathscr{A}$ be the inverse image of $\mathscr{A}$ under the morphism $f\colon Y \to X$. Since $(f^{-1}\mathscr{A})_y$ is isomorphic to $\mathscr{A}_x$ and the former is an ultra-ring by the above, so is therefore the latter. □

In particular, any stalk over an atomless Boolean scheme is an ultra-ring!

http://www.springer.com/978-3-642-13367-1

The Use of Ultraproducts in Commutative Algebra
Schoutens, H.
2010, X, 210 p., Softcover
ISBN: 978-3-642-13367-1