

ESSENTIALLY ALGEBRAIC
THEORIES AND LOCALIZATIONS IN
TOPOSES AND ABELIAN
CATEGORIES

A THESIS SUBMITTED TO THE UNIVERSITY OF MANCHESTER
FOR THE DEGREE OF DOCTOR OF PHILOSOPHY
IN THE FACULTY OF ENGINEERING AND PHYSICAL SCIENCES

2012

Philip Bridge
School of Mathematics

Contents

Abstract	4
Declaration	5
Copyright Statement	6
Acknowledgements	8
1 Introduction	9
2 Localizations and geometric logic	13
2.1 Locally presentable categories	13
2.2 Essentially algebraic theories	23
2.3 Localization for toposes	31
2.4 Localization in abelian categories	37
2.5 Regular theories in the language of R -modules	42
2.6 Models of a geometric theory	48
2.7 Classifying toposes	52
3 Sheaves as essentially algebraic objects	56
3.1 Sheaves as essentially algebraic objects	56
3.2 Locally coherent and finitely presented toposes	57
3.3 Modules over a sheaf of rings	66
4 Krull-Gabriel Dimension	77
4.1 Supercompactness and regularity	77

4.2	Modular lattices and m-dimension	88
4.3	Krull-Gabriel dimension for abelian categories	93
4.4	Krull-Gabriel dimension for toposes	96
4.5	The Ziegler frame	100
4.6	A Ziegler spectrum for abelian groups	103
4.7	An example with a cyclic group	108
5	Flat Covers in Algebraic Categories	113
5.1	Strongly flat models of a theory	113
5.2	The Flat Cover Conjecture	120
5.3	Flat covers in presheaf categories	121
	Bibliography	131

The University of Manchester

Philip Bridge

Doctor of Philosophy

Essentially Algebraic Theories and Localizations in Toposes and Abelian Categories

January 27, 2012

The main theme of this thesis is the parallel between results in topos theory and in the theory of additive functor categories.

In chapter 2, we provide a general overview of the topics used in the rest of the thesis. Locally finitely presentable categories are introduced, and their expression as essentially algebraic categories is explained. The theory of localization for toposes and abelian categories is introduced, and it is shown how these localizations correspond to theories in appropriate logics.

In chapter 3, we look at conditions under which the category of modules for a ring object R in a topos \mathcal{E} is locally finitely presented, or locally coherent. We show that if \mathcal{E} is locally finitely presented, then the category of modules is also; however, we show that far stronger conditions are required for the category of modules to be locally coherent.

In chapter 4, we show that the Krull-Gabriel dimension of a locally coherent abelian category \mathcal{C} is equal to the socle length of the lattice of regular localizations of \mathcal{C} . This is used to make an analogous definition of Krull-Gabriel dimension for regular toposes, and the value of this dimension is calculated for the classifying topos of the theory of G -sets, where G is a cyclic group admitting no elements of square order.

In chapter 5, we introduce a notion of strong flatness for algebraic categories (in the sense studied by Adamek, Rosicky and Vitale in [3]). We show that for a monoid M of finite geometric type, or more generally a small category \mathcal{C} with the corresponding condition, the category of M -acts, or more generally the category of set-valued functors on \mathcal{C} , has strongly flat covers.

Declaration

No portion of the work referred to in this thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright Statement

- i The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the "Copyright") and s/he has given the University of Manchester certain rights to use such Copyright, including for administrative purposes.
- ii Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made **only** in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.
- iii The ownership of certain Copyright, patents, designs, trade marks and other intellectual property (the "Intellectual Property") and any reproductions of copyright works in the thesis, for example graphs and tables ("Reproductions"), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.
- iv Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see <http://www.campus.manchester.ac.uk/medialibrary/policies/intellectualproperty.pdf>), in any relevant Thesis restriction declarations deposited in the University Library, The University Library's regulations

(see <http://www.manchester.ac.uk/library/aboutus/regulations>) and in the University's policy on presentation of Theses.

Acknowledgements

I would like to thank my supervisor, Mike Prest, for all his help and support during the research that led to this thesis.

I would also like to thank Harold Simmons, Victoria Gould, Alex Bailey and John MacQuarrie, and my office mates Simon, Richard, Lorna, Erik and Wemedh, for many helpful discussions and distractions.

Finally, I would like to thank my girlfriend, Eiko, for all her support and patience over the last four years.

Chapter 1

Introduction

The main theme of this thesis is to investigate connections between Grothendieck abelian categories and Grothendieck toposes. Both types of category are examples of locally presentable categories. The similarity between these two fields has been known for a long time, with abelian functor categories playing a role analogous to an ‘additive version’ of the classifying topos associated with the theory of modules over a ring. In particular, taking logical quotients of this theory corresponds to looking at localizations of the functor category in the abelian case, or the classifying topos in the non-abelian case. This theory is outlined in the book by Prest [23] and the thesis of Caramello [11]. In chapter 2 of this thesis we introduce locally presentable categories and explain how to interpret them as categories of models of an essentially algebraic theory. This material is from [1]. We then introduce toposes and abelian categories, and explain the theory of their localizations, before showing that these correspond to theories in the appropriate geometric logics. This material is drawn together from the books [20], [23], [17] and [11], and none of the results are new, though the proof of 2.28 is new. It uses the ideas in the proof of the Duality Theorem in [11], and is quite different from the argument given in [23].

The main idea in chapter 3 is to see sheaves over a site as models of an essentially algebraic theory. This can be used to classify those toposes which are locally finitely presented, generated and coherent (3.2, 3.4, 3.7). These characterizations are all known, but we believe this method of proving them is new. We then look at a

question considered in the paper [24], which asked, given a ring object R in a locally finitely presented topos \mathcal{E} , when the category of modules over R is itself locally finitely presented. It is not difficult to see using the theorems about essentially algebraic theories from [1] that the category of R -modules is also locally finitely presented (Theorem 3.11). The remainder of the chapter is spent examining when the category of R -modules might be locally coherent. Using an argument of Prest and Ralph, we show that if R is a finitely presented sheaf of rings on a space such that RU is a coherent ring for every open set U in the space, then the category of modules over R is locally coherent (Theorem 3.14). On the other hand, we show that for any space with infinitely many open set, we will be able to choose a sheaf of rings R such that all the rings RU are coherent, but the category of modules is not locally coherent (Theorem 3.16).

In chapter 4, we examine the idea of regular localizations and Krull-Gabriel dimension. We show that if \mathbb{T} is a Cartesian theory, then the regular objects in the classifying topos for \mathbb{T} can be considered as pairs of regular formulas ϕ/θ , where θ defines an equivalence relation on ϕ . This is a ‘regular’ non-additive analogue of Burke’s result, [23, 10.2.30], which states for a ring R , the finitely presented functors $\text{mod-}R \rightarrow \mathbf{Ab}$ are given by pairs of pp-formulas ϕ/ψ , where $\psi(M) \subseteq \phi(M)$ for every R -module M . A ‘coherent’ version of this result was proved in [26], and our proof is almost identical; it is the simple observation that at each stage of Rajani’s proof of this result, if all the formulas we start with are regular, than all the formulas constructed during the proof will be regular also. We give an ‘internal’ definition of localizations of the classifying topos corresponding to regular theories (that is, one not mentioning the theory explicitly, but only the objects of the topos) - a localization is regular if and only if the inclusion functor preserves epimorphic families of monomorphisms. Finally, we characterize locally regular toposes using an argument similar to that used for the coherent case in chapter 3.

In the next part of the chapter, we look for a good definition of Krull-Gabriel dimension for toposes. Krull-Gabriel dimension is an invariant of locally coherent abelian categories, which to some extent measures the complexity of the categories’

structure. It is defined using successive finite-type localizations of the category. Our objective is to come up with a good analogous definition for toposes. For Grothendieck abelian categories, regular and finite-type localization coincide, but this is not the case for toposes. We will find that the notion of dimension we get using regular localizations is a far better invariant than we would reach using finite-type localizations. We show that the Krull-Gabriel dimension of a locally coherent abelian category \mathcal{C} is equal to the socle length of the lattice of finite-type/regular localizations of \mathcal{C} (Corollary 4.29). This allows us to make the same definition for a locally regular topos \mathcal{E} . Using this definition, we calculate the dimension for the classifying topos for the theory of \mathbb{Z}_n -sets, for a cyclic group \mathbb{Z}_n , and show that the dimension will be equal to the number of distinct prime factors of n (Proposition 4.42). To do this we construct a ‘Ziegler spectrum’ for this category - that is, a set of objects in the category sufficient to classify regularly definable subcategories (section 4.6). This construction is based on the one given in the paper [28], but corrects an error made in that paper.

In chapter 5 we consider the Flat Cover Conjecture for toposes. This result, stating that every object in a Grothendieck abelian category has a flat cover, was first proved for categories of modules in [6], and more recently for all locally finitely presented abelian categories [12, Corollary 3.3], [31, Theorem 2]. We extend the concept of flatness to cover not just toposes but arbitrary algebraic categories, that is, those categories which can be expressed as the categories of finite product-preserving functors $\mathcal{T} \rightarrow \mathbf{Sets}$, where \mathcal{T} is an arbitrary category with finite products. This coincides with the property of strong flatness that has previously been investigated by semigroup theorists, e.g., in [21]. In particular, we show that for a monoid M of finite geometric type (in fact, for a small category \mathcal{C} with the corresponding property) the category of functors $M \rightarrow \mathbf{Sets}$ (or the category of functors $\mathcal{C} \rightarrow \mathbf{Sets}$ if we consider the categorical version) has strongly flat covers (Theorem 5.21).

We will assume the reader is familiar with the basics of category theory (limits and colimits, adjoint functors, the Yoneda lemma, etc.) as outlined in, for example, [7]. The books [1] and [20] will be referred to frequently, and arguments that are given in full in these books will be outlined only very briefly in this thesis. A familiarity

with model theory will be useful in reading the thesis, but not essential.

Categories will generally be denoted using calligraphic capital letters (e.g., \mathcal{C} , \mathcal{D} , ...); tuples will generally be denoted with bold-faced lower case letters (e.g., \mathbf{x} , \mathbf{y} , ...). For categories \mathcal{C} and \mathcal{D} , the category of functors $\mathcal{C} \rightarrow \mathcal{D}$ with morphisms the natural transformations will generally be denoted $(\mathcal{C}, \mathcal{D})$. This contrasts to the notation $\text{Hom}_{\mathcal{C}}(C, D)$, will be used to denote the set of morphisms $C \rightarrow D$ between two objects in \mathcal{C} . Where the category in question is obvious from the context, the subscript \mathcal{C} will sometimes be omitted. An italicized word or phrase (e.g., *free realization*) will indicate that it is being defined, unless given as part of a definition, in which case the word or phrase being defined will be in normal typeface. Lemmas, definitions and theorems share the same counter, which is reset at the beginning of each chapter, so Lemma 4.3 refers to the third result in chapter 4.

Chapter 2

Localizations and geometric logic

2.1 Locally presentable categories

We begin by describing the various presentability conditions on categories. The notions of finite presentability and coherence introduced in this section are from the books by Adamek and Rosicky [1], Johnstone [17] and Prest [23].

A partially ordered set (I, \leq) is said to be *directed* if it is non-empty and any two elements of I have a least upper bound. Let \mathcal{C} be a category. A *directed system* in \mathcal{C} is a functor $D : \mathcal{D} \rightarrow \mathcal{C}$ where \mathcal{D} is a directed partially ordered set, considered as a category (that is, we consider the elements of the partial order to be objects, and for $x, y \in \mathcal{D}$, there is a unique arrow $x \rightarrow y$ if and only if $x \leq y$ in the order). A *directed colimit* in \mathcal{C} is a colimit over a directed diagram. We denote the directed colimit over a diagram D with the notation $\varinjlim D$. Directed colimits are often called direct limits, especially by algebraists. An object C in \mathcal{C} is said to be *finitely presentable* if the functor

$$\mathrm{Hom}(C, -) : \mathcal{C} \rightarrow \mathbf{Sets}$$

commutes with directed colimits. Equivalently, given a directed system

$$\{D_i \xrightarrow{d_{ij}} D_j \mid i \leq j \in (I, \leq)\}$$

where (I, \leq) is some directed poset, and given a colimit cocone

$$\{D_i \xrightarrow{d_i} L\}$$

we have that any map $f : C \rightarrow L$ factors through the cocone, that is, $f = d_i f'$ for some $i \in I$ and some $f' : C \rightarrow D_i$, and this factorization is essentially unique, in the sense that if $g : C \rightarrow D_j$ is some other map with $d_j g = f$, then for some $k \geq i, j$, we have $d_{ik} f' = d_{jk} g$.

A set of objects \mathcal{G} is said to *generate* \mathcal{C} if for any pair of arrows $f \neq g : A \rightarrow B$ in \mathcal{C} , there is some map $x : G \rightarrow A$ with $G \in \mathcal{G}$ and $fx \neq gx$. Equivalently, if \mathcal{C} has coproducts, for every object C there is an epimorphism

$$e : \coprod_i G_i \rightarrow C$$

where the objects G_i are all in \mathcal{G} .

An epimorphism is $e : E \rightarrow C$ said to be *strong* if given any commutative square

$$\begin{array}{ccc} A & \xrightarrow{e} & B \\ f \downarrow & & \downarrow g \\ E & \xrightarrow{m} & C \end{array}$$

such that m is a monomorphism, there is a map $d : B \rightarrow E$ such that $md = g$ and $f = de$. In categories with pushouts, this is equivalent to stating for e is *extremal*, that is, it does not factor through any proper subobject of C .

If \mathcal{C} has coproducts, a generating set \mathcal{G} is said to *strongly generate* \mathcal{C} if for every object C there is a strong epimorphism $e : \coprod_i G_i \rightarrow C$ as above. This equivalent to the condition that whenever $s : S \rightarrow A$ is a proper monomorphism in \mathcal{C} , there is a map $x : G \rightarrow A$ with $G \in \mathcal{G}$, not admitting a factorization through s .

The category \mathcal{C} is *locally finitely presentable* if it is cocomplete and has a strong generating set of finitely presentable objects.

Analogous to the above, we say an object C in \mathcal{C} is *finitely generated* if the

representable functor $\text{Hom}(C, -)$ commutes with directed colimits of diagrams where all the maps d_{ij} are monics. Such a diagram is called a *directed union*. The category \mathcal{C} is called *locally finitely generated* if it is cocomplete, co-wellpowered and has a strong generating set of finitely generated objects (note that one can prove locally finitely presentable categories are co-wellpowered using the definition above (e.g., [1, 1.58, 2.49]); it is an open question whether this condition is necessary in the definition of a locally finitely generated category).

Finally, an object C in \mathcal{C} is *coherent* if it is finitely generated, and for any pullback diagram of the form

$$\begin{array}{ccc} B \times_C B' & \longrightarrow & B' \\ \downarrow & & \downarrow \\ B & \longrightarrow & C \end{array}$$

in which B and B' are finitely generated, we have that $B \times_C B'$ is also finitely generated. A category is *locally coherent* if it is cocomplete and has a strong generating set of coherent objects.

Example. 1. Let S be a set (of sorts). An *S -sorted set* is a collection of sets $(X_s)_{s \in S}$ indexed by S ; a morphism of S -sorted sets $f : X \rightarrow Y$ consists of a family of functions $f_s : X_s \rightarrow Y_s$, indexed by S .

An *S -sorted signature of (finitary) algebras* Σ consists of a collection of function symbols σ , together with an arity function assigning to each function symbol σ an ordered sequence (s_1, \dots, s_n) from S denoting the domain and an element $s \in S$ denoting the codomain. We write this information as: $\sigma : s_1 \times \dots \times s_n \rightarrow s$.

An *algebra* A of the signature Σ consists of an S -sorted set $|A| = (A_s)_{s \in S}$ together with functions $\sigma_A : A_{s_1} \times \dots \times A_{s_n} \rightarrow A_s$ for each function symbol $\sigma : s_1 \times \dots \times s_n \rightarrow s$. For nullary function symbols (i.e., those for which $n = 0$) σ_A will denote simply an element of A_s .

A homomorphism of algebras $A \rightarrow B$ consists of an S -sorted map $|A| \rightarrow |B|$, preserving the operations in Σ , in the sense that for any function symbol $\sigma : s_1 \times \dots \times s_n \rightarrow s$, and any set of elements $x_1 \in A_{s_1}, \dots, x_n \in A_{s_n}$, the equation

$f_s(\sigma_A(x_1, \dots, x_n)) = \sigma_B(f_{s_1}(x_1), \dots, f_{s_n}(x_n))$ will hold.

The category of algebras over Σ is a locally finitely presentable category.

2. Let Σ be a many-sorted signature of algebras, and let X be an S -sorted set (of variables). The set of terms over X is defined as follows: each variable $x \in X_s$ is a term of sort S , and, if $\sigma : s_1 \times \dots \times s_n \rightarrow s$ is an operation symbol, and t_i is a term of sort s_i for $i = 1, \dots, n$, then $\sigma(t_1, \dots, t_n)$ is a term of sort s . We write $T_\Sigma(X)$ for the set of all terms over the S -sorted set X . An *equation* in the variables X is a pair of terms in $T_\Sigma(X)$, (t_1, t_2) of the same sort; we write this as $t_1 = t_2$. An algebra A of the signature is said to *satisfy* the equation $t_1 = t_2$ if the equation holds for each interpretation of the variables X in Σ (i.e., for each S -sorted function $X \rightarrow A$, the elements of A corresponding to the terms t_1 and t_2 will be equal).

An *equational theory* (Σ, E) over a set of sorts S consists of an S -sorted signature of algebras together with a set of equations E (in some ‘standard’ set of variables X). A *model* of the theory consists of an algebra A of the signature Σ in which all the equations in E are satisfied.

Let (Σ, E) be an equational theory over a set of sorts S . The category of models of (Σ, E) is locally finitely presented [1, 3.7]. The free algebras on the sorts of S are finitely presentable, and form a strong generating set for the category.

Examples of categories of models of equational theories include groups, rings, monoids, R -modules (where R is a given ring), Lie algebras over a given field k , etc. Note that the precise signature we choose for each theory can affect whether the theory is equational for that signature. For example, the theory of groups is an equational theory over the single-sorted signature (m, e, i) , where m is binary, i is unary and e is a constant; considered as a special class of monoid (i.e., a theory over the signature (m, e)), groups are not described by an equational theory.

3. An *S -sorted relational signature* is a collection of relation symbols Σ , together

with an arity function assigning to each relation symbol ρ of Σ an ordered sequence $(s_1, \dots, s_n) \in S^n$, denoting the domain of the relation ρ . A *relational structure* A of type Σ consists of an S -sorted set $|A|$ together with, for each $\rho \in \Sigma$ of arity (s_1, \dots, s_n) , a relation $\rho_A \subseteq A_{s_1} \times \dots \times A_{s_n}$. A homomorphism of relational structures $f : A \rightarrow B$ consists of an S -sorted map $f : |A| \rightarrow |B|$ preserving the relations, i.e., so that for each relation $\rho \in \Sigma$ of arity (s_1, \dots, s_n) , we have that $(x_1, \dots, x_n) \in \rho_A$ implies $(f_{s_1}(x_1), \dots, f_{s_n}(x_n)) \in \rho_B$. For an S -sorted relational structure Σ , the category of Σ -structures will be locally finitely presentable.

4. An *S -sorted signature* Σ consists of a disjoint union $\Sigma_{\text{fun}} \cup \Sigma_{\text{rel}}$, where Σ_{fun} is an S -sorted signature of algebras, and Σ_{rel} is an S -sorted relational signature. A Σ -structure is an S -sorted set A together with operations and relations that make A both an algebra of Σ_{fun} and a relational structure of type Σ_{rel} ; a morphism of S -sorted structures is an S -sorted map $f : |A| \rightarrow |B|$ which is both a morphism of algebras and a morphism of relational structures. The category of Σ -structures and their morphisms will again be a locally finitely presentable category.

Locally finitely presentable categories are interesting from the point of view of model theory because it is possible to see the objects in them as models of a multi-sorted theory, with the finitely presentable objects in the category being the sorts, and the maps to the object from a finitely presentable object as being the elements of that sort in the model.

Lemma 2.1. ([27, 2.1]) *Let \mathcal{C} be a locally finitely generated category, with a generating set \mathcal{G} of finitely generated objects. An object C of \mathcal{C} is finitely generated if there is a strong epimorphism*

$$\coprod_{i=1}^n G_i \xrightarrow{e} C$$

where for each i , $G_i \in \mathcal{G}$.

Lemma 2.2. ([27, 2.2]) *Let \mathcal{C} be a locally finitely presentable category, with a*

generating set \mathcal{G} of finitely presentable objects. An object C of \mathcal{C} is finitely presentable if and only if there is a coequaliser diagram

$$\coprod_{j=1}^m H_j \rightrightarrows \coprod_{i=1}^n G_i \longrightarrow C$$

where for every i and j , $G_i, H_j \in \mathcal{G}$.

Remark. In fact, it follows from the proof in [27] that for *any* strong epimorphism $e : \coprod_{i=1}^n G_i \rightarrow C$, we can construct a coequaliser diagram of this form.

We recall the definition of an exact category, in the sense of Barr, which can be found in, for example [8, ch.2] or [16, A1.3]. There is some inconsistency in these sources as to whether an exact category is required to have all finite limits, since only specific limits are required to prove most of the theorems about them; in practice most of the examples that are studied do have finite limits. In this thesis we will assume that they do, for simplicity.

A category \mathcal{C} is said to be *regular* if it has finite limits, every kernel pair has a coequaliser, and the pullback of a strong epimorphism along any morphism is again a strong epimorphism.

Now let $(a, b) : R \rightrightarrows C$ be a parallel pair of morphisms in a category \mathcal{C} with finite limits.

1. We say (a, b) is a *relation* on C if $(a, b) : R \rightarrow C \times C$ is monic.
2. We say (a, b) is *reflexive* if there exists $r : C \rightarrow R$ with $ar = br = 1_C$.
3. We say (a, b) is *symmetric* if there exists $s : R \rightarrow R$ with $as = b$ and $bs = a$.
4. We say (a, b) is *transitive* if there exists $t : P \rightarrow R$, such that $at = ap$ and $bt = bq$, where P is the pullback

$$\begin{array}{ccc} P & \xrightarrow{q} & R \\ p \downarrow & & \downarrow a \\ R & \xrightarrow{b} & C \end{array} .$$

5. We say that (a, b) is an *equivalence relation* if it has all four of the above properties.

6. We say that the equivalence relation (a, b) is *effective* if it has a coequaliser $q : C \rightarrow Q$, and (a, b) is the kernel pair of q .

A category \mathcal{C} is said to be *exact* if it is regular and every equivalence relation in \mathcal{C} is effective.

Remark. This form of exactness is also called ‘exact in the sense of Barr’ to distinguish it from another notion of categorical exactness concerning additive categories, due to Quillen. The notion above is the only notion of exactness we will be using in this thesis.

Corollary 2.3. *In an exact locally finitely presentable category, an object C is finitely presentable if there is a strong epimorphism $e : B \rightarrow C$ with B finitely presentable, where the underlying object for the kernel pair of e is finitely generated.*

Proof. Let $f, f' : K \rightarrow B$ be the kernel pair. If K is finitely generated there is a strong epimorphism $g : L \rightarrow K$ with L finitely presentable. Because g is a strong epimorphism, the diagram

$$L \begin{array}{c} \xrightarrow{fg} \\ \xrightarrow{f'g} \end{array} B \xrightarrow{e} C$$

is a coequaliser diagram. □

Corollary 2.4. *Let \mathcal{C} be an exact locally finitely presentable category. Then every coherent object is finitely presentable.*

Proof. Let \mathcal{G} be a generating set of finitely presentable objects for \mathcal{C} .

Let C be a coherent object of \mathcal{C} . There is a strong epimorphism $e : B \rightarrow C$, with B finitely presentable. Since C is coherent, the underlying object for the kernel pair of e is finitely generated. Thus C is finitely presentable. □

Lemma 2.5. *In any category \mathcal{C} , if C is coherent and $s : S \rightarrow C$ is a finitely generated subobject of C , then S is coherent also.*

Proof. Suppose we have a pullback diagram of the form

$$\begin{array}{ccc} B \times_S B' & \longrightarrow & B' \\ \downarrow & & \downarrow b' \\ B & \xrightarrow{b} & S \end{array}$$

For any pair of maps $x : X \rightarrow B$, $x' : X \rightarrow B'$ we have $bx = b'x'$ if and only if $sbx = sb'x'$. Thus the diagram

$$\begin{array}{ccc} B \times_S B' & \longrightarrow & B' \\ \downarrow & & \downarrow sb' \\ B & \xrightarrow{sb} & C \end{array}$$

is a pullback also, and $B \times_S B'$ is finitely generated. □

Remark. In a locally coherent abelian category, coherent objects are characterised by the condition that their finitely generated subobjects are finitely presented [23, 2.3.15]; however this condition is not equivalent to coherence in general. If it was, then the next result would show that finitely presentable and coherent objects coincided, and in locally coherent abelian categories, this is the case [23, E.1.18].

Theorem 2.6. *Let \mathcal{C} be an exact locally coherent category, in which coherent objects are closed under finite coproducts. Then any finitely generated subobject of a finitely presentable object is finitely presentable.*

Proof. Let C be a finitely presentable object, with a finitely generated subobject $s : S \rightarrow C$. There is a strong epimorphism $a : A \rightarrow S$, where A is coherent. Since C is finitely presentable, there is a strong epimorphism $b : B \rightarrow C$ where B is coherent, A can be taken to be a subobject of B , and the object underlying the kernel pair of b is finitely generated (to see that A is a subobject of B , choose an arbitrary coherent object admitting C as a quotient, and take the coproduct of this with A). Consider the following diagram:

$$\begin{array}{ccccc}
 K_1 & \xrightarrow{f_1} & A & \xrightarrow{a} & S \\
 & \xrightarrow{f_2} & \downarrow s' & & \downarrow s \\
 K_2 & \xrightarrow{g_1} & B & \xrightarrow{b} & C \\
 & \xrightarrow{g_2} & & &
 \end{array}$$

Here (f_1, f_2) is the kernel pair of a , and (g_1, g_2) is the kernel pair of b . According to our assumptions, B and A are coherent, and K_2 is finitely generated.

Since $af_1 = af_2$, we have $bs'f_1 = saf_1 = saf_2 = bs'f_2$, and so there is a unique map $t : K_1 \rightarrow K_2$ with $s'f_1 = g_1t$, $s'f_2 = g_2t$.

We claim that the cone

$$\begin{array}{ccc}
 K_1 & \xrightarrow{f_1} & A \\
 & \xrightarrow{f_2} & \\
 \downarrow t & & \\
 K_2 & &
 \end{array}$$

is a limit over the diagram

$$\begin{array}{ccc}
 & A & \\
 & \downarrow s' & \swarrow s' \\
 K_2 & \xrightarrow{g_1} & B \\
 & \xrightarrow{g_2} &
 \end{array}$$

To see this, suppose we have an object X and arrows $x_1, x_2 : X \rightarrow A$, $x_3 : X \rightarrow K_2$, satisfying $s'x_1 = s'x_2 = g_1x_3 = g_2x_3$.

Since $s'x_1 = s'x_2$, we have $bs'x_1 = bs'x_2$, and so $sax_1 = sax_2$. Since s is monic, this means $ax_1 = ax_2$. Since (f_1, f_2) is the kernel pair of a , there is a map $x : X \rightarrow K_1$ with $f_1x = x_1$, $f_2x = x_2$. It suffices to prove that $tx = x_3$. But this is true because both of these maps are factorizations of the map $s'x_1 = s'x_2$ through the kernel pair (g_1, g_2) , so they are equal by the uniqueness of this factorization.

Thus K_1 is the limit of the diagram described. But this limit can be constructed by taking three pullbacks, one after the other - that is, K_1 is isomorphic to the composition of pullbacks $((K_2 \times_{(g_1, g_2)} K_2) \times_B A) \times_B A$. Since K_2 and A are each finitely generated, and B is coherent, it follows that these pullbacks, and therefore K_1 , are finitely generated. Thus S is finitely presentable. \square

The notions of finite presentability introduced in this section can be generalized to infinite cardinals. Let λ be a regular cardinal. A partially ordered set (I, \leq) is said

to be λ -directed if every subset $S \subseteq I$ of size less than λ has a least upper bound. In a category \mathcal{C} , a λ -directed diagram is a functor $D : \mathcal{D} \rightarrow \mathcal{C}$ where the category \mathcal{D} is a λ -directed partially ordered set. A λ -directed colimit is a colimit over a λ -directed diagram. An object C in \mathcal{C} is said to be λ -presented if the functor $\text{Hom}(C, -)$ commutes with λ -directed colimits, and λ -generated if $\text{Hom}(C, -)$ commutes with λ -directed colimits where all the maps in the colimit cocone are monic; such a diagram is called a λ -directed union. The category \mathcal{C} is said to be *locally λ -presented* if it is cocomplete and has a strong generating set of λ -presented objects, and it is said to be *locally λ -generated* if it is cocomplete, cowellpowered, and has a strong generating set of λ -generated objects. We say \mathcal{C} is *locally presentable* if there is a regular cardinal λ such that \mathcal{C} is locally λ -presented.

Let \mathcal{C} be a category, and let \mathcal{A} be a reflective subcategory, with reflection functor $r : \mathcal{C} \rightarrow \mathcal{A}$ and inclusion functor $i : \mathcal{A} \rightarrow \mathcal{C}$. This reflection is of λ -type if i preserves λ -directed colimits (equivalently, \mathcal{A} is closed under λ -directed colimits). We say the reflection is of *finite type* if it is of \aleph_0 type.

In this thesis, we will often use the next result, from [1, 1.39], which shows that λ -type reflections preserve local λ -presentability. We include a proof of this result, since we will refer back to it a lot during this thesis.

Theorem 2.7. ([1, 1.39]) *Let \mathcal{C} be a locally λ -presentable category and let $r : \mathcal{C} \rightarrow \mathcal{A}$ be a λ -type reflection functor onto a full subcategory $\mathcal{A} \subseteq \mathcal{C}$. Then the subcategory \mathcal{A} is also locally λ -presentable.*

If \mathcal{C} is locally λ -generated, and $r : \mathcal{C} \rightarrow \mathcal{A}$ is a reflection functor such that the inclusion functor i preserves λ -directed unions, then the subcategory \mathcal{A} is also locally λ -generated.

Proof. First, we show that if an object C in \mathcal{C} is λ -presented (respectively λ -generated), then the object rC is λ -presented (respectively, λ -generated) in \mathcal{A} .

To see this, let

$$\{D_i \xrightarrow{d_i} L\}$$

be a λ -directed colimit (respectively, a λ -directed union) in \mathcal{A} , and suppose we have

a map $f : rC \rightarrow L$. Now the cocone is a colimit in \mathcal{C} also (since the inclusion functor preserves it), so the map $f\eta_C : C \rightarrow rC \rightarrow L$ factors through the cocone, say as $f\eta_C = d_i g$ for some map $g : C \rightarrow D_i$. But by the reflection property, g factors through the map η_C , say as $g = g'\eta_C$ for some $g' : rC \rightarrow D_i$. We have $d_i g'\eta_C = d_i g = f\eta_C$, so by uniqueness of the factorisation, $d_i g' = f$, and g' is the required factorisation.

Now let $\mathcal{G} = \{G_i\}_{i \in I}$ be a strong generating set for \mathcal{C} . We claim that $r\mathcal{G} = \{rG_i\}_{i \in I}$ is a strong generating set for \mathcal{A} . To see this set has the separating property, suppose we have a parallel pair of morphisms $f \neq g : A \rightarrow A'$ in \mathcal{A} . There is a map $x : G_i \rightarrow A$ with $fx \neq gx$ for some $G_i \in \mathcal{G}$ (since \mathcal{G} is a generating set for \mathcal{C}). This map factors uniquely through rG_i , to give $x' : rG_i \rightarrow A$, with $fx' \neq gx'$.

To see that the generator is strong, suppose we have a proper monomorphism $s : S \rightarrow A$ in \mathcal{A} . This is a proper monomorphism in \mathcal{C} also, so there is a map $x : G_i \rightarrow A$ such that x does not factor through s . The factorisation $x' : rG_i \rightarrow A$ cannot factor through s also. It follows that $r\mathcal{G}$ is a strong generating set for \mathcal{A} , and in particular if \mathcal{G} is taken to be the set of finitely presentable objects in \mathcal{C} , we see that \mathcal{A} is locally λ -presented (respectively, locally λ -generated). \square

2.2 Essentially algebraic theories

Locally presentable categories can be characterized as categories of models of essentially algebraic theories. This gives a nice way of thinking about the objects in these categories. In this section, we will introduce essentially algebraic theories, and give an explicit description of an essentially algebraic theory associated with any given locally presentable category.

Recall the following from [1, 3.34].

Definition 2.8. 1. *An essentially algebraic theory is given by a quadruple*

$$\Gamma = (\Sigma, E, \Sigma_t, Def)$$

Here Σ is a many sorted signature of algebras, over some set of sorts S .

Σ_t is a subset of Σ , denoting the set of function symbols we intend to view as total. We write Σ_p for the set $\Sigma - \Sigma_t$; these function symbols are to be interpreted as being partial. As usual, we view constant symbols as function symbols defined over the empty set of sorts.

The set E consists of equations over Σ between terms in variables x_i , where each x_i has a sort $s_i \in S$.

Finally Def is a function assigning to each partial function symbol $\sigma : \prod_{i \in I} s_i \rightarrow s$ a collection of Σ_t -equations in variables $x_i \in s_i$, ($i \in I$). These equations are taken to define the domain of definition for σ .

2. We say that Γ is λ -ary for a regular cardinal λ if each function symbol in Σ takes fewer than λ arguments, and each $Def(\sigma)$ contains fewer than λ equations.
3. By a model of Γ , we mean a partial Σ -algebra A such that A satisfies all equations of E , the total functions are everywhere defined, and a partial function $\sigma \in \Sigma_p$ is defined for a tuple $\mathbf{a} \in A^{\alpha(\sigma)}$ if and only if the tuple \mathbf{a} satisfies all the equations in $Def(\sigma)$.

It is well-known [1, 3.36] that the categories of models of λ -ary essentially algebraic theories are precisely the locally λ -presentable categories. We will present a proof of this fact which we hope will make clear a sense in which an essentially algebraic object can be seen as being generated by a collection of its elements, in a similar manner to the way an ordinary algebraic object is.

Now assume we are given an essentially algebraic theory $\Gamma = (\Sigma, E, \Sigma_t, Def)$ and a collection of variables $x_1 \in S_1, \dots, x_n \in S_n$, we construct terms in Γ over these variables as follows:

1. each variable x_i is a term of sort S_i .
2. given a total operation $f : S_1 \times \dots \times S_n \rightarrow S$, and terms $t_1 \in S_1, \dots, t_n \in S_n$, $f(t_1, \dots, t_n)$ is a term of sort S .

3. if $\sigma : S_1 \times \dots \times S_n \rightarrow S$ is a partial operation and $t_1 \in S_1, \dots, t_n \in S_n$ are terms such that the equations $\text{Def}(\sigma)(t_1, \dots, t_n)$ hold in every model of (Σ, E) , then $\sigma(t_1, \dots, t_n)$ is a term of sort S .

Remark. The stipulation that the equations $\text{Def}(\sigma)(t_1, \dots, t_n)$ hold in every model of (Σ, E) in step 3 above could be replaced, via a suitable Completeness Theorem, with the assertion that the axioms E admit a deduction of each of the equations in $\text{Def}(\sigma)(t_1, \dots, t_n)$, either in classical first order logic or via the regular deductive calculus described in section 2.5.

In a normal algebraic theory (that is, one where all the operations are assumed to be total), a presentation of an object is given by a collection of generators \mathbf{x} and a conjunction of equations $\phi(\mathbf{x})$ which we assert is satisfied by \mathbf{x} . When we allow partial operations we add the complication that the equations which go together to form ϕ should include only terms over \mathbf{x} which are well-defined. Consequently, we make the following definition: a *presentation* consists of a collection of variables $\mathbf{x} = \{x_i \in S_i\}_{i \in I}$ and a totally ordered collection of equations $R(\mathbf{x}) = \{\phi_i(\mathbf{x}) \mid i \in I\}$, where (I, \leq) is some total order, and each ϕ_i is an equation in terms over \mathbf{x} . This has the restriction that if the partial operation σ is used to form a term in $\phi_i(\mathbf{x})$, then the arguments for σ satisfy the equations $\text{Def}(\sigma)$ in every model of $E \cup \{\phi_j(\mathbf{x}) \mid j < i\}$.

We define terms over a presentation as follows:

1. each variable x_i is a term of sort S_i .
2. given a total operation $f : S_1 \times \dots \times S_n \rightarrow S$, and terms $t_1 \in S_1, \dots, t_n \in S_n$, $f(t_1, \dots, t_n)$ is a term of sort S .
3. if $\sigma : S_1 \times \dots \times S_n \rightarrow S$ is a partial operation and $t_1 \in S_1, \dots, t_n \in S_n$ are terms such that the equations $\text{Def}(\sigma)(t_1, \dots, t_n)$ are satisfied in every model of $(\Sigma, E \cup R(\mathbf{x}))$, then $\sigma(t_1, \dots, t_n)$ is a term of sort S .

The collection of terms over a presentation is itself a model $\langle \mathbf{x} \mid R(\mathbf{x}) \rangle$ of Γ , with the obvious operations. This is universal, in the sense that for any other Γ -model Y , and any tuple $\mathbf{y} \in Y$ satisfying $R(\mathbf{y})$, there is a unique map $\tilde{\mathbf{y}} : \langle \mathbf{x} \mid R(\mathbf{x}) \rangle \rightarrow Y$

mapping the variables \mathbf{x} to \mathbf{y} . The definition of $\tilde{\mathbf{y}}$ is given on the terms over the presentation by induction.

Furthermore, any Γ -model C admits such a presentation - we can take \mathbf{x} to be all the elements of C , and $R(\mathbf{x})$ to be all the equations holding between them. We will of course usually be able to find smaller presentations than this.

Colimits can be determined in terms of these presentations. Let A and B be models of Γ , with presentations $\langle \mathbf{x}, R(\mathbf{x}) \rangle, \langle \mathbf{y}, S(\mathbf{y}) \rangle$. Then the coproduct $A \coprod B$ is the object with presentation $\langle (\mathbf{x}, \mathbf{y}), R(\mathbf{x}) \cup S(\mathbf{y}) \rangle$. Given a Γ -model C and mappings $f : A \rightarrow C, g : B \rightarrow C$, we can find corresponding tuples $\mathbf{a} \in R(C)$ and $\mathbf{b} \in S(C)$. The tuple $(\mathbf{a}, \mathbf{b}) \in R \cup S(C)$ corresponds to the coproduct factorization.

Similarly, given a parallel pair of morphisms $f, g : A \rightarrow B$, for every $a \in A$, the elements $f(a)$ and $g(a)$ can be expressed as terms $t_a^f(\mathbf{y}), t_a^g(\mathbf{y})$ over the generators \mathbf{y} of B . A map $h : B \rightarrow C$ with $hf = hg$, corresponds to a tuple $\mathbf{c} \in S(C)$ with the additional property that for each $a \in A, t_a^f(\mathbf{c}) = t_a^g(\mathbf{c})$ (this is the condition that $hf = hg$). In fact it suffices to require this just for the generators \mathbf{x} of A . Thus the coequaliser of f and g admits the presentation $\langle \mathbf{y} \mid S(\mathbf{y}) \cup \{t_x^f(\mathbf{y}) = t_x^g(\mathbf{y})\}_{x \in \mathbf{x}} \rangle$.

In particular, given a Γ -model B and a presentation $\langle \mathbf{y} \mid S(\mathbf{y}) \rangle$ of B , each equation τ in S is of the form $t_1^\tau(\mathbf{y}) = t_2^\tau(\mathbf{y})$. Furthermore t_1^τ and t_2^τ have the same sort, X_τ , say. Let B_1 be the free Γ -model on generators $\mathbf{x} = \{x_\tau \in X_\tau\}$, and B_2 the free model on generators \mathbf{y} . There is a pair of maps $f, g : B_1 \rightarrow B_2$ defined by $f : x_\tau \mapsto t_1^\tau(\mathbf{y}), g : x_\tau \mapsto t_2^\tau(\mathbf{y})$. Then B is the coequaliser of the maps f and g .

We summarize this information.

Lemma 2.9. *Every model C of an essentially algebraic theory Γ admits a presentation $\langle \mathbf{x} \mid R(\mathbf{x}) \rangle$, and C can be expressed as the coequaliser of a diagram*

$$F \rightrightarrows G \longrightarrow C$$

where F and G are free models of Γ . The number of generators of G is bounded by the cardinality of \mathbf{x} , and the number of generators of F is bounded by the cardinality of $R(\mathbf{x})$.

For algebras over a signature Σ , the notions of finitely presented and finitely generated correspond to the usual notions we can define using generators and relations, this is proved in [1, 3.10]. We seek now to prove an analogous result for essentially algebraic objects. Let $\Gamma = (\Sigma, E, \Sigma_t, \text{Def})$ be an essentially algebraic theory. Let X be a model of Γ . We say that a tuple of elements $\mathbf{x} \in X$ *generates* X if every element of X can be written as a term over the elements \mathbf{x} , using the term forming operations as described above.

Lemma 2.10. *Let $\Gamma = (\Sigma, E, \Sigma_t, \text{Def})$ be a λ -ary essentially algebraic theory. Then the forgetful functor $U : \mathbf{Mod}(\Gamma) \rightarrow \mathbf{Mod}(\Sigma_t, E_t)$ preserves λ -directed colimits.*

Proof. It is sufficient to show that if (D_i, \leq) is a λ -directed system of Γ -structures, then the colimit of the underlying (Σ_t, E_t) -structures is also a Γ -structure. To see this, let

$$D_i \xrightarrow{d_i} L$$

be the colimit cocone. Suppose we have some tuple \mathbf{x} of elements of L , that satisfy the equations $\text{Def}(\sigma)$, for some partial operation σ . Then by the definition of a λ -directed colimit, there is some $i \in I$ and some $\mathbf{x}' \in D_i$ such that $D_i \models \text{Def}(\sigma)(\mathbf{x}')$ and $d_i(\mathbf{x}') = \mathbf{x}$. The well-defined term $d_i(\sigma(\mathbf{x}'))$ then gives us our definition for $\sigma(\mathbf{x})$. \square

Lemma 2.11. *(cf [1, 3.11], [27, 3.16(a)]) Let Γ be a λ -ary essentially algebraic theory. A model X is a λ -generated object in $\mathbf{Mod}(\Gamma)$ if and only if it has a generating set of size less than λ .*

Proof. Suppose X is λ -generated. For each set of elements $S \subseteq X$ with size less than λ , let \bar{S} be the substructure of X generated by S (this is the set of all elements which can be written as terms over the elements of S). X can be written as the union of all the \bar{S} ; therefore $X = \bar{S}$ for some S .

To prove the converse, suppose X has a generating set $\{x_i\}_{i \in S}$ of size $|S|$ less than λ , and let X be the λ -directed union of a collection of subobjects $X = \bigcup U_j$. Then each x_i can be written as a term over the elements of the U_j , and since the terms are λ -small, this term can only involve elements of less than λ of the objects U_j . Since

there are fewer than λ elements x_i , the set of U_j 's needed to write all the x_i as terms over the elements of the U_j 's is also of cardinality less than λ . \square

Lemma 2.12. (cf [1, 3.12], [27, 3.16(b)]) *The λ -presented objects in $\text{Mod}(\Gamma)$ are precisely those which have a λ -small presentation.*

Proof. We show that a free model F of Γ on a single generator of sort X is λ -presented. This follows from the fact that λ -directed colimits in $\text{Mod}(\Gamma)$ are calculated as in (Σ_t, E_t) . Thus, given a directed colimit cocone

$$\{D_i \xrightarrow{d_i} L \mid i \in I\}$$

over a directed system $\{d_{ij} : D_i \rightarrow D_j \mid i \leq j \in (I, \leq)\}$ for some directed poset (I, \leq) , a map $F \rightarrow L$ corresponds to an element $x \in L$ of sort X . But since the directed colimit is the same as that for the underlying (Σ_t, E_t) structures, there is some $i \in I$ and some $x' \in D_i$ with $d_i(x') = x$. Thus the map $F \rightarrow D_i$ defined by x' is an appropriate factorization through the cocone.

Having established that a free object on a single generator is λ -presented, the result now follows from 2.9 and 2.2. \square

Essentially algebraic theories characterise locally presentable categories; that is, a category \mathcal{C} is locally presentable if and only if it is the category of models for an essentially algebraic theory. This is proved in [1, 3.36]. We will give a different proof of this, which will describe explicitly an essentially algebraic theory associated with a given locally presentable category.

To prove this result, we introduce the following concept, from [1, 1.42]. For a small category \mathcal{A} and a regular cardinal λ , denote by $\mathbf{Cont}_\lambda \mathcal{A}$ the category of all functors $\mathcal{A} \rightarrow \mathbf{Sets}$ preserving all λ -small limits in \mathcal{A} .

Theorem 2.13. ([1, 1.46]) *If \mathcal{C} is a locally finitely presentable category, and \mathcal{A} is the subcategory of finitely presentable objects in \mathcal{C} , then \mathcal{C} is equivalent to $\mathbf{Cont}_\lambda \mathcal{A}^{op}$.*

Theorem 2.14. *A category \mathcal{C} is locally λ -presentable if and only if it is equivalent to the category of models of a λ -ary essentially algebraic theory $(\Sigma, E, \Sigma_t, \text{Def})$.*

Furthermore, this is a reflective subcategory of the category of models of the equational theory (Σ_t, E_t) , where E_t is the subset of E containing those equations not using any of the function symbols from Σ_p and the inclusion functor preserves λ -directed colimits.

Proof. We have already proved the last part.

Suppose we are given a λ -ary essentially algebraic theory $\Gamma = (\Sigma, E, \Sigma_t, \text{Def})$. Let C be a model of (Σ_t, E_t) . The reflection of C is just the Γ -structure given by the presentation $\langle \mathbf{c}, R(\mathbf{c}) \rangle$, where the variables in \mathbf{c} are the elements of C , and $R(\mathbf{c})$ is the set of in Σ_t holding for the elements of C (with an arbitrary ordering).

If \mathcal{C} is a locally λ -presentable category and \mathcal{A} is a reflective subcategory closed under λ -directed colimits, then \mathcal{A} is also locally λ -presentable. The reflections of the λ -presentable objects in \mathcal{C} are λ -presentable in \mathcal{A} , and form a strong generating set [1, 1.3].

To show that every locally λ -presentable category can be represented this way, let \mathcal{C} be a locally λ -presented category, with \mathcal{A} the category of λ -presented objects in \mathcal{C} .

Define a λ -ary essentially algebraic theory Γ as follows. The total part of Γ is just the category \mathcal{A}^{op} , with equations those holding in \mathcal{A}^{op} .

For each λ -small diagram $\mathcal{D} = \{D_i \xrightarrow{f_{ij}^k} D_j\}$ in \mathcal{A} , with colimit cocone $\{D_i \xrightarrow{d_i} L\}$, the object L will be in \mathcal{A} , since λ -presentable objects are closed under λ -small colimits. Define partial operations $\sigma_{\mathcal{D}} : \prod_{i \in I} D_i \rightarrow L$, where $\text{Def}(\sigma)$ is the set of equations $f_{ij}^k(x_j) = x_i$ for each arrow $f_{ij}^k : D_i \rightarrow D_j$ (note we can have $D_i = D_j$ for $i \neq j$). Add equations to our theory stating that for any tuple $\mathbf{x} = \{x_i \in D_i\}_{i \in I} \in \prod_{i \in I} D_i$, then $d_i \sigma_{\mathcal{D}}(\mathbf{x}) = x_i$, for each $i \in I$.

The category of models for Γ is the category of presheaves on \mathcal{A} which preserve the λ -small limits existing in \mathcal{A}^{op} ; that is, the category $\mathbf{Cont}_{\lambda} \mathcal{A}^{\text{op}}$ of λ -continuous set-valued functors on \mathcal{A}^{op} . By Theorem 2.13, this is equivalent to \mathcal{C} . This proves that Γ is an essentially algebraic theory whose category of models is equivalent to \mathcal{C} . □

Remark. The Yoneda embedding gives us a way to see the objects of \mathcal{C} as models of

the theory Γ in the obvious way: let C be an object of \mathcal{C} . Then for each object A in \mathcal{A} , the set of elements of \mathcal{C} of that sort is the collection of maps $A \rightarrow C$; the functions in \mathcal{A} act on this set by precomposition. For each diagram \mathcal{D} , the map $\sigma_{\mathcal{D}}$ sends a compatible cocone over the diagram with codomain \mathcal{C} to the factorisation through the colimit.

It is clear from this that the free Γ -models in the given presentation are precisely the λ -presented objects in \mathcal{C} .

Lemma 2.15. *Let Γ be an essentially algebraic theory $(\Sigma, E, \Sigma_t, Def)$ such that every function symbol in Σ is finitary. Then the category of Γ -models is locally finitely generated.*

Proof. It suffices to show that for an essentially algebraic theory of the above form, if (I, \leq) is a directed poset and $\{d_{ij} : D_i \rightarrow D_j \mid i \leq j \in I\}$ is a directed union of Γ -structures, then the colimit of the underlying (Σ_t, E_t) -structures is also a Γ -structure.

Write $\{d_i : D_i \rightarrow D \mid i \in I\}$ for the colimit cocone. Note that each d_i is also a monic map. In a locally presentable category, the monic maps are precisely the injective maps, so we can consider the D_i 's to be essentially algebraic substructures of D .

Now let σ be a partial operation, with domain of definition given by $Def(\sigma)$. Since σ is a finitary operation, the set of equations $Def(\sigma)$ uses only finitely many variables. Let \mathbf{d} be a tuple in D such that $D \models Def(\sigma)(\mathbf{d})$. Since the D_i 's cover D , each element d_k from \mathbf{d} occurs as an element of D_i for some i . Since the D_i 's occur as a directed system, we can find some D_j containing the whole tuple \mathbf{d} . Then we define $\sigma(\mathbf{d})$ to be $d_j(\sigma^{D_j}(\mathbf{d}))$. \square

We conjecture that locally finitely generated categories are characterized by essentially algebraic theories of this form, but we do not have a proof of this.

Definition 2.16. [1, 2.27] *Let \mathcal{C} be a category.*

- A monomorphism $m : A \rightarrow B$ is said to be pure if for each commutative square

$$\begin{array}{ccc} A' & \xrightarrow{f} & B' \\ u \downarrow & & \downarrow v \\ A & \xrightarrow{m} & B \end{array}$$

where A' and B' are finitely presentable, u factorizes through f .

- An epimorphism $e : A \rightarrow B$ is said to be pure if for each morphism $f : C \rightarrow B$ with C finitely presentable, there is a factorization of f through e .

Lemma 2.17. *Let \mathcal{C} be a locally finitely presentable category, and let $r : \mathcal{C} \rightarrow \mathcal{A}$ be a finite type reflection. Then r preserves pure monics and pure epics.*

Proof. By [1, 2.30] and [2, 4], pure monics are directed colimits of split monics and pure epics are directed colimits of split epics. The claim then follows from the fact that finite type reflections preserve directed colimits. \square

2.3 Localization for toposes

In this section, we introduce the ideas of Grothendieck topologies, sheaves and toposes. The material in this section is covered in [20, III.2,4].

Throughout this section, let \mathcal{C} be a small category. A *presheaf* on \mathcal{C} is a functor $P : \mathcal{C}^{op} \rightarrow \mathbf{Sets}$.

For an object C in \mathcal{C} , a *sieve* S on C is a subfunctor of the representable functor $\mathrm{Hom}(-, C)$. This is equivalent to a collection of morphisms with codomain C , such that given morphisms f, g with $f \in S$ and $\mathrm{cod}(g) = \mathrm{dom}(f)$, then $fg \in S$.

If S is a sieve on C and $f : D \rightarrow C$ is a map in \mathcal{C} , then the collection

$$f^*(S) = \{g : E \rightarrow D \mid fg \in S\}$$

is a sieve on D , and the diagram

$$\begin{array}{ccc} f^*(S) & \longrightarrow & \text{Hom}(-, D) \\ [g \mapsto fg] \downarrow & & \downarrow [g \mapsto fg] \\ S & \longrightarrow & \text{Hom}(-, C) \end{array}$$

is a pullback diagram in $\mathbf{Sets}^{C^{\text{op}}}$.

A *Grothendieck topology* J on \mathcal{C} assigns to each object C in \mathcal{C} a collection JC of sieves on C , called *covering sieves*, such that:

1. for each C , $\text{Hom}(-, C) \in JC$.
2. if $S \in JC$ and $f : D \rightarrow C$ is any map, then $f^*(S) \in JD$.
3. if $S \in JC$ and R is a sieve on C such that $h^*(R) \in J(\text{dom}(h))$ for all $h \in S$, then $R \in JC$.

A *site* (\mathcal{C}, J) consists of a small category \mathcal{C} equipped with a Grothendieck topology J .

To specify a sieve S , it suffices to specify a collection of morphisms S' which generates the sieve, in the sense that for every morphism $f : A \rightarrow C$ in S , there is a morphism $g : B \rightarrow C$ in S' , such that $f = gh$ for some $h : A \rightarrow B$. For this reason, a collection of morphisms with common codomain is sometimes called a *presieve*. Grothendieck topologies can be defined in terms of presieves: a *basis* for a Grothendieck topology on \mathcal{C} is a function K assigning to each object C a set of presieves on C , called *covering families*, such that

1. if $f : C' \rightarrow C$ is an isomorphism, then $\{f : C' \rightarrow C\} \in KC$.
2. if $\{f_i : C_i \rightarrow C \mid i \in I\} \in KC$ and $g : D \rightarrow C$ is any morphism then there exists some covering family $\{h_j : D_j \rightarrow D \mid j \in J\} \in KD$ such that for each $j \in J$, gh_j factors through some map f_i .
3. if $\{f_i : C_i \rightarrow C \mid i \in I\} \in KC$ and for each $i \in I$ we have a covering family

$$\{g_{ij} : D_{ij} \rightarrow C_i \mid j \in J_i\} \in KC_i,$$

then

$$\{f_i g_{ij} : D_{ij} \rightarrow C \mid i \in I, j \in J_i\} \in KC.$$

Note that if \mathcal{C} has pullbacks, then in condition (2) we may require that the morphisms h_j are precisely the pullbacks of the morphisms f_i along g .

Given a basis K for a Grothendieck topology, we obtain a topology J by stating that the covering sieves in J are precisely those which contain a presieve in K . Every topology J has a basis; given a topology J , we define KC for each object C to be the set of all families of morphisms $F = \{f_i : C_i \rightarrow C\}$ with codomain C such that there exists a sieve $S \in JC$ for which the family of morphisms F is *cofinal*, that is, such that every morphism in S factors through some morphism in F .

Example. Let X be a topological space. Write $\text{Op}(X)$ for the lattice of open sets of X ; we regard $\text{Op}(X)$ as a small category, with the inclusion maps as morphisms. There is a Grothendieck topology on $\text{Op}(X)$: for an open set $U \in \text{Op}(X)$ a family of maps $\{U_i \rightarrow U \mid i \in I\}$ is a covering family if $U = \bigcup_{i \in I} U_i$.

For the rest of this section, assume we have fixed a Grothendieck topology J on \mathcal{C} .

A presheaf $F : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Sets}$ is said to be a *sheaf* for the topology J if whenever we have a covering sieve S of an object C and a diagram in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ of the form

$$\begin{array}{ccc} S & \xrightarrow{\mathbf{x}} & F \\ \downarrow i & & \\ \text{Hom}(-, C) & & \end{array}$$

there is a unique natural transformation $x : \text{Hom}(-, C) \rightarrow F$ such that $xi = \mathbf{x}$.

If F is a sheaf on the site (\mathcal{C}, J) , where the category \mathcal{C} is the lattice $\text{Op}(X)$ for some topological space X , and J is the Grothendieck topology described above, we say F is a sheaf on X .

To explain our choice of notation: a natural transformation $S \rightarrow F$ consists of identifying, for each $f : D \rightarrow C$ in S , an element $x_f \in FD$, in such a way that whenever $f \in S$ and g is a morphism such that $\text{cod}(g) = \text{dom}(f)$, then $x_{fg} = Fg(x_f)$.

We call such a collection of elements $\mathbf{x} = \{x_f \mid f \in S\}$ a *matching family* for S of elements of F .

An *amalgamation* for a matching family is an element $x \in FC$ such that for each $f \in S$, $x_f = Ff(x)$. Identifying x with the corresponding natural transformation $\text{Hom}(-, C) \rightarrow F$ via the Yoneda lemma, we see that the sheaf condition is the precise statement that every matching family has a unique amalgamation.

We write $\mathbf{Sh}(\mathcal{C}, J)$ for the category of sheaves on the site (\mathcal{C}, J) , considered as a full subcategory of $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ (that is, we take as morphisms the natural transformations between the functors).

A (*Grothendieck*) *topos* is a category equivalent to $\mathbf{Sh}(\mathcal{C}, J)$ for some (not necessarily unique) site (\mathcal{C}, J) . In particular, we can define the *trivial topology* on \mathcal{C} to be the topology J for which the only covering sieve on each object C is the functor $\text{Hom}(-, C)$. We can see straight away that every presheaf is a sheaf for this topology, and so the presheaf category $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ is a Grothendieck topos.

There is another notion of topos occurring in the literature, referred to as an *elementary topos*. This approach defines a topos to be a category with certain of the ‘nice’ properties held by a category of sheaves of sets (and in particular, a Grothendieck topos is an elementary topos also). We will not make use of this concept, and throughout the rest of the thesis the word ‘topos’ shall be used to mean a Grothendieck topos, but the reader should be aware this usage is not uniform throughout the literature.

Lemma 2.18. ([20, IV.1.2]) *Let \mathcal{E} be a topos. Any morphism $f : A \rightarrow B$ in \mathcal{E} which is both monic and epic is an isomorphism.*

A category \mathcal{C} with the property that any morphism which is both epic and monic is an isomorphism is called a *balanced* category. Toposes and abelian categories are examples of balanced categories.

The following definition is from [7]:

Definition 2.19. *Let \mathcal{C} be any category. A localization of \mathcal{C} is a reflection functor $l : \mathcal{C} \rightarrow \mathcal{A}$ that preserves finite limits.*

There is a localization functor $a : \mathbf{Sets}^{C^{\text{op}}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$. We call the functor a the *associated sheaf* functor. We give an outline of the construction of this functor now; the details of many of the claims made here are omitted, though they may be checked rather easily. The argument is given in detail in [20, III.5], and we refer the reader there for a complete version of the construction.

The associated sheaf functor is defined by two applications of the ‘plus-functor’ $P \mapsto P^+$. For an arbitrary presheaf P , this is given by

$$P^+C = \lim_{\rightarrow S \in JC} \text{Match}(S, P),$$

that is, each element of P^+C is an equivalence class of matching families for covers of C where two matching families $\{x_f\}_{f \in S}$ and $\{y_g\}_{g \in T}$ are equivalent if there is a common refinement of S and T on which they agree, that is, if there exists $U \subseteq S \cap T$ such that for all $f \in U$, $x_f = y_f$.

Now suppose $h : C' \rightarrow C$ is a morphism in \mathcal{C} . For each dense sieve S on C , we define a function $\text{Match}(S, P) \rightarrow \text{Match}(h^*S, P)$ by

$$\{x_f \mid f \in S\} \mapsto \{x_{hf'} \mid f' \in h^*S\}.$$

This induces a map $P^+C \rightarrow P^+C'$; defining P^+h to be this map, it can be shown that P^+ is a presheaf. Each map $\phi : P \rightarrow Q$ of presheaves induces a map $\phi^+ : P^+ \rightarrow Q^+$ of presheaves by taking a matching family $S \twoheadrightarrow P$ to the composite $S \twoheadrightarrow P \xrightarrow{\phi} Q$.

There is a map of presheaves $\chi : P \rightarrow P^+$ defined by sending each $x \in PC$ to the equivalence class of $\{Pf(x) \mid f \in \text{Hom}(-, C)\}$.

If F is a sheaf and P is a presheaf, then any map $\phi : P \rightarrow F$ of presheaves factors uniquely through η ; we write $\phi = \tilde{\phi}\chi$, as in the diagram

$$\begin{array}{ccc} P & \xrightarrow{\chi} & P^+ \\ & \searrow \phi & \downarrow \tilde{\phi} \\ & & F \end{array}$$

Moreover, for every presheaf P , $(P^+)^+$ is a sheaf. We define $aP = P^{++}$; this gives

us a functor $a : \mathbf{Sets}^{\mathcal{C}^{\text{op}}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$ which is a left adjoint to the inclusion functor $i : \mathbf{Sh}(\mathcal{C}, J) \rightarrow \mathbf{Sets}^{\mathcal{C}^{\text{op}}}$. We denote the unit of the adjunction by η ; for each presheaf P , the map $\eta_P : P \rightarrow aP$ has the property that any map $P \rightarrow F$ with F a sheaf factors uniquely through η_P .

We have shown that for a small category \mathcal{C} , every Grothendieck topology J on \mathcal{C} gives rise to a localization $a_J : \mathbf{Sets}^{\mathcal{C}^{\text{op}}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$. In fact, a converse result holds, in the following sense: every localization $a : \mathbf{Sets}^{\mathcal{C}^{\text{op}}} \rightarrow \mathcal{E}$ corresponds to a Grothendieck topology J_a on \mathcal{C} . For an object C in \mathcal{C} , the J_a -covering sieves on C are those sieves S where the inclusion morphism $S \rightarrow \text{Hom}(-, C)$ is mapped to an isomorphism by a . Thus there is a bijective correspondance between the Grothendieck topologies on \mathcal{C} and the localizations of \mathcal{C} .

Moreover, the collection of Grothendieck topologies on \mathcal{C} forms a complete lattice - it is clear that Grothendieck topologies are closed under arbitrary intersections; the join of a collection of Grothendieck topologies $\{J_i\}_{i \in I}$ on \mathcal{C} is the intersection of all the topologies J that contain every one of the J_i (i.e., $J_i \leq J$ for all $i \in I$).

For an arbitrary Grothendieck topos $\mathbf{Sh}(\mathcal{C}, J)$, the localizations correspond to those Grothendieck topologies on \mathcal{C} which contain J . These Grothendieck topologies are closed under the lattice operations. Thus we see that for any Grothendieck topos \mathcal{E} , the localizations of \mathcal{E} form a complete lattice.

We conclude this section by mentioning the following general notion of a morphism of toposes.

Definition 2.20. *Let \mathcal{E} and \mathcal{F} be toposes. A geometric morphism $f : \mathcal{E} \rightarrow \mathcal{F}$ consists of a pair of functors $f^* : \mathcal{F} \rightarrow \mathcal{E}$ and $f_* : \mathcal{E} \rightarrow \mathcal{F}$, such that f^* preserves finite limits and is left adjoint to f_* . The functor f^* is called the inverse image of the morphism, and the functor f_* is called the direct image of the morphism.*

Let (\mathcal{C}, J) be some site. The functors $a : \mathbf{Sets}^{\mathcal{C}^{\text{op}}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$ and $i : \mathbf{Sh}(\mathcal{C}, J) \rightarrow \mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ form a geometric morphism, such that the inclusion i is full and faithful. We call such a geometric morphism an *embedding*.

2.4 Localization in abelian categories

In this section, we summarize the theory of localization for Grothendieck abelian categories. An abelian category is said to be *Grothendieck* if it cocomplete, has a generating set of objects, and finite limits commute with directed colimits (see, e.g. [23, p. 707]). In particular, a locally finitely presentable abelian category is Grothendieck. A Grothendieck abelian category is analogous to a Grothendieck topos, and similarly, locally finitely presentable abelian categories can be seen as parallels to locally finite presentable toposes, and so on. In the case of toposes, we described a localization of the topos by listing some morphisms which were to be inverted. In the abelian case, we can instead denote objects, which are to be made equal to zero; this is equivalent to inverting those morphisms whose kernels and cokernels are amongst these objects.

Let \mathcal{C} be an abelian category. An abelian subcategory $\mathcal{B} \subseteq \mathcal{C}$ is said to be *closed under extensions* if for an exact sequence

$$0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0$$

with A, C objects in \mathcal{B} , the object B is in \mathcal{B} also.

A *Serre subcategory* of an abelian category \mathcal{C} is a non-empty, full subcategory \mathcal{S} closed under extensions, subobjects and quotients.

For an additive functor $F : \mathcal{C} \rightarrow \mathcal{C}'$ between abelian categories, the *kernel* of F is the full subcategory of \mathcal{C} whose objects are precisely those objects C with $F(C) \cong 0$ in \mathcal{C}' .

The Serre subcategories of \mathcal{C} are the kernels of exact functors with domain \mathcal{C} , in much the same way that ideals in ring theory are the kernels of ring morphisms.

Theorem 2.21. ([22, 4.3.3], [23, 11.1.40]) *Given a Serre subcategory \mathcal{S} of an abelian category \mathcal{C} , there is a quotient category \mathcal{C}/\mathcal{S} , which is also abelian, and there is an exact functor $Q_{\mathcal{S}} : \mathcal{C} \rightarrow \mathcal{C}/\mathcal{S}$, such that $\ker(Q_{\mathcal{S}}) = \mathcal{S}$ and $Q_{\mathcal{S}}$ is universal with this property, in the sense that for any other exact functor $F : \mathcal{C} \rightarrow \mathcal{C}'$ with*

$\mathcal{S} \subseteq \ker(F)$, there is a unique factorization of F through $Q_{\mathcal{S}}$ via an exact and faithful functor.

It is clear from the definition that Serre subcategories are closed under arbitrary intersections. Given a collection of objects \mathcal{A} in \mathcal{C} , the Serre subcategory $\mathcal{S}_{\mathcal{A}}$ is the intersection of all the Serre subcategories containing \mathcal{A} . This is described explicitly by the next result, which is a slight generalization of [23, 13.1.2]; the proof is almost identical.

Lemma 2.22. *Let \mathcal{A} be a collection of objects in \mathcal{C} closed under quotients and subobjects; then the Serre subcategory $\mathcal{S}_{\mathcal{A}}$ generated by \mathcal{A} consists precisely of objects F in \mathcal{C} such that there is a finite composition series*

$$0 = F_0 < F_1 \dots < F_n = F$$

such that for each $i \leq n$, the quotient F_i/F_{i-1} is in \mathcal{A} .

Proof. It is clear that every object F with this property is in $\mathcal{S}_{\mathcal{A}}$. It suffices to show that the collection of such objects is closed under quotients, subobjects and extensions.

Let F be an object with a composition series as described, and let S be a subobject of F . Let $S_n = S$, and for each $i = n - 1, n - 2, \dots, 1, 0$, define S_i to be the pullback of the maps $S_i \rightarrow F_i$ and $F_{i-1} \rightarrow F_i$. We get a composition series for S :

$$\begin{array}{ccccccc} 0 = S_0 & \longrightarrow & S_1 & \longrightarrow & \dots & \longrightarrow & S_n = S \\ \downarrow & & \downarrow & & & & \downarrow \\ 0 = F_0 & \longrightarrow & F_1 & \longrightarrow & \dots & \longrightarrow & F_n = F \end{array}$$

Now for each $i = 0, 1, \dots, n - 1$, we have a canonical map $S_{i+1}/S_i \rightarrow F_{i+1}/F_i$. This map is monic (see for example, [8, 1.10.2]). Since $F_{i+1}/F_i \in \mathcal{A}$, it follows that S_{i+1}/S_i is also in \mathcal{A} .

Now suppose we have a quotient map $q : F \rightarrow Q$. For each $i = 0, \dots, n$, we define

Q_i to be the image of the map

$$F_i \longrightarrow F \xrightarrow{q} Q.$$

This gives us a composition series for Q :

$$\begin{array}{ccccccc} 0 = F_0 & \longrightarrow & F_1 & \longrightarrow & \dots & \longrightarrow & F_n = F \\ \downarrow & & \downarrow & & & & \downarrow \\ 0 = Q_0 & \longrightarrow & Q_1 & \longrightarrow & \dots & \longrightarrow & Q_n = Q \end{array}$$

Again, for each $i = 1, \dots, n$, we have a canonical map $e_i : F_i/F_{i-1} \rightarrow Q_i/Q_{i-1}$, which we claim is an epimorphism. For each i , the map $F_i \rightarrow Q_i \rightarrow Q_i/Q_{i-1}$ is an epimorphism, and factors through e_i , so e_i must be an epimorphism.

Each Q_i/Q_{i-1} must be in \mathcal{A} , since it is a quotient of F_i/F_{i-1} .

We have shown that the collection of objects with the property described is closed under subobjects and quotients. It remains to show that it is closed under extensions.

Suppose we have an exact sequence in \mathcal{A} of the form

$$0 \longrightarrow F \xrightarrow{f} G \xrightarrow{g} H \longrightarrow 0$$

where both F and H admit composition series of the form described.

$$0 = F_0 \longrightarrow F_1 \longrightarrow \dots \longrightarrow F_m = F$$

$$0 = H_0 \longrightarrow H_1 \longrightarrow \dots \longrightarrow H_n = H$$

Define $G_n = G$. For each $i = n - 1, n - 2, \dots, 1, 0$, define G_i to be the pullback of the diagram

$$\begin{array}{ccc} & & G_{i+1} \\ & & \downarrow \\ H_i & \longrightarrow & H_{i+1} \end{array}$$

We get a composition series for G :

$$\begin{array}{ccccccc}
 G_0 & \longrightarrow & G_1 & \longrightarrow & \dots & \longrightarrow & G_n = G \\
 \downarrow g_0 & & \downarrow g_1 & & & & \downarrow g_n = g \\
 0 = H_0 & \longrightarrow & H_1 & \longrightarrow & \dots & \longrightarrow & H_n = H
 \end{array}$$

Now the diagram

$$\begin{array}{ccc}
 G_0 & \longrightarrow & G \\
 \downarrow & & \downarrow g \\
 0 & \longrightarrow & H
 \end{array}$$

is a pullback diagram. However, the pullback of this diagram will be the kernel of g , which by assumption was F . Thus G_0 is isomorphic to F . In particular, F is contained in each of the objects G_i , and in fact it is the kernel of each morphism $g_i : G_i \rightarrow H_i$. For if $x : X \rightarrow G_i$ satisfies $g_i x = 0$, then there is a (unique) factorisation of $g_i x$ through the 0 object, and thus of g_i through F .

We have shown that each object H_i is of the form H_i/F . This allows us to apply the Noether isomorphism theorems (for example see [8, 1.10.6]): each quotient G_i/G_{i-1} is isomorphic to $(G_i/F)/(G_{i-1}/F)$, and therefore to H_i/H_{i-1} . Thus, the composition series

$$0 = F_0 \longrightarrow F_1 \longrightarrow \dots \quad F_m = F = G_0 \longrightarrow G_1 \longrightarrow \dots \quad G_n = G$$

is of the required form, and G has the required property. This concludes the proof. \square

Let \mathcal{C} be a Grothendieck abelian category, with generating set \mathcal{G} . When considering Grothendieck abelian categories, we look for functors which preserve the directed colimits in the category — for exact functors, this is equivalent to demanding that (possibly infinite) coproducts are preserved. So we will consider those subcategories of \mathcal{C} which are kernels of exact functors preserving directed colimits.

Thus, a *hereditary torsion class* in \mathcal{C} is defined as a subcategory \mathcal{T} closed under extensions, subobjects, quotients and arbitrary direct sums (i.e., coproducts).

Theorem 2.23. ([23, 11.1.5]) *Let \mathcal{C} be a Grothendieck abelian category, and let \mathcal{T} be a hereditary torsion class on \mathcal{C} . There is a Grothendieck abelian subcategory $\mathcal{C}_{\mathcal{T}}$*

in \mathcal{C} , and a localization functor $r_{\mathcal{T}} : \mathcal{C} \rightarrow \mathcal{C}_{\mathcal{T}}$ such that for each $T \in \mathcal{T}$, $rT = 0$, and $r_{\mathcal{T}}$ is universal with this property, in the sense that for any exact functor $F : \mathcal{C} \rightarrow \mathcal{D}$ admitting a right adjoint, with $\mathcal{T} \subseteq \ker(f)$, there is a unique factorization of F through $r_{\mathcal{T}}$.

We will in particular be interested in the theory of localization for locally coherent abelian categories. We mention some properties of locally coherent abelian categories now. If \mathcal{C} is a locally finitely presentable abelian category, it is locally coherent if and only if every finitely presentable object in \mathcal{C} is coherent [23, E.1.18], or equivalently, if the full subcategory of coherent objects in \mathcal{C} is an abelian subcategory [23, E.1.19].

We describe the localizations that preserve the property of local coherence. A hereditary torsion class is said to be of *finite type* if the inclusion functor $i : \mathcal{C}_{\mathcal{T}} \rightarrow \mathcal{C}$ preserves directed unions.

Theorem 2.24. ([23, 11.1.34]) *If \mathcal{C} is a locally coherent abelian category, and \mathcal{T} is a finite-type hereditary torsion class on \mathcal{C} , then the category $\mathcal{C}_{\mathcal{T}}$ is also locally coherent, and the coherent objects in $\mathcal{C}_{\mathcal{T}}$ are precisely the images of the coherent objects in \mathcal{C} under the localization.*

Let \mathcal{C} be a locally coherent category with a finite-type hereditary torsion class \mathcal{T} as described above. Suppose the subcategory of coherent objects in \mathcal{C} is denoted \mathcal{A} . Then $\mathcal{T} \cap \mathcal{A}$ is a Serre subcategory of \mathcal{A} . Moreover, the restriction of $Q_{\mathcal{T}}$ to \mathcal{A} is the quotient functor $Q_{\mathcal{T} \cap \mathcal{A}} : \mathcal{A} \rightarrow \mathcal{A}/(\mathcal{T} \cap \mathcal{A})$ [23, 11.1.42]. Conversely, if \mathcal{S} is a Serre subcategory of \mathcal{A} , then if $\mathcal{T}_{\mathcal{S}}$ is the closure of \mathcal{S} in \mathcal{C} under directed colimits, the category $\mathcal{T}_{\mathcal{S}}$ is a finite-type hereditary torsion class, and the functor $Q_{\mathcal{T}_{\mathcal{S}}}$ is an extension of the functor $Q_{\mathcal{S}}$ on \mathcal{A} .

This sets up a 1-1 correspondence between Serre subcategories of \mathcal{A} and finite-type hereditary torsion classes on \mathcal{C} [23, 12.4.1].

Analogously to the case with toposes, we note that Serre subcategories of a small abelian category, or alternatively finite-type hereditary torsion classes in a locally coherent abelian category, are closed under arbitrary intersections. Thus, they form a complete lattice.

2.5 Regular theories in the language of R -modules

Let Σ be a many-sorted signature of algebras. A *regular formula* (also called a *positive primitive formula*) in the language Σ is a formula $\phi(\mathbf{x})$ of the form $\exists \mathbf{y} \wedge_{i=1}^n \psi_i(\mathbf{x}, \mathbf{y})$ where each formula $\psi_i(\mathbf{x}, \mathbf{y})$ is an atomic formula in the language, i.e., an equation in variables (\mathbf{x}, \mathbf{y}) .

A *regular sequent* is a statement of the form

$$\phi(\mathbf{x}) \vdash \psi(\mathbf{x}),$$

where ϕ and ψ are regular formulas whose free variables are included in the variables \mathbf{x} . The intention is that in the above statement ϕ and ψ are interpreted as formulas in the variables \mathbf{x} , regardless of whether all of these variables actually occur explicitly in the formula. Given a Σ -structure X , the sequent $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$ is said to *hold* in X if given any choice of the variables \mathbf{x} from the elements of X , the formula $\psi(\mathbf{x})$ holds if the formula $\phi(\mathbf{x})$ does.

A *regular theory* \mathbb{T} over a signature Σ is a set of regular sequents over Σ , whose elements are called the *axioms* of \mathbb{T} .

Regular sequents admit rules of deduction as detailed below. These rules of deduction are described in more detail in [17, D1.3] (and we have deliberately labeled our rules for existential quantification to be consistent with that book). Note though that in that book, the free variables occurring in a sequent are written as subscripts to the logic gate; we use the notation described above in order to refer to the variables \mathbf{x} over which a formula is taken without reference to a sequent.

(a) The structural rules consist of the identity axiom

$$\phi(\mathbf{x}) \vdash \phi(\mathbf{x}),$$

the substitution rule

$$\frac{\phi(\mathbf{x}) \vdash \psi(\mathbf{x})}{\phi'(\mathbf{y}) \vdash \psi'(\mathbf{y})},$$

where the formulas $\phi'(\mathbf{y})$, $\psi'(\mathbf{y})$ are obtained by substituting the variables \mathbf{y} for the variables \mathbf{x} in $\phi(\mathbf{x})$ and $\psi(\mathbf{x})$, and the cut rule

$$\frac{\phi(\mathbf{x}) \vdash \psi(\mathbf{x}) \quad \psi(\mathbf{x}) \vdash \chi(\mathbf{x})}{\phi(\mathbf{x}) \vdash \chi(\mathbf{x})}.$$

(b) The equality rules consist of the axioms

$$\top \vdash (x = x)$$

and

$$(x = y) \wedge \phi(\mathbf{x}) \vdash \phi'(\mathbf{y}).$$

(c) The rules for finite conjunction are the axioms

$$\phi(\mathbf{x}) \vdash \top, \quad (\phi(\mathbf{x}) \wedge \psi(\mathbf{x})) \vdash \phi(\mathbf{x}), \quad (\phi(\mathbf{x}) \wedge \psi(\mathbf{x})) \vdash \psi(\mathbf{x})$$

and the rule

$$\frac{\phi(\mathbf{x}) \vdash \psi(\mathbf{x}) \quad \phi(\mathbf{x}) \vdash \chi(\mathbf{x})}{\phi(\mathbf{x}) \vdash (\psi \wedge \chi)(\mathbf{x})}.$$

(f) The rules for existential quantification consists of the double rule (i.e., either statement may be deduced from the other)

$$\frac{\phi(\mathbf{x}, y) \vdash \psi(\mathbf{x}, y)}{(\exists y)\phi(\mathbf{x}, y) \vdash \psi(\mathbf{x})}.$$

Here we assume the variable y does not occur freely in ψ , so it does make sense to talk of the formula $\psi(\mathbf{x})$.

The following is proved in [17, D1.5.4].

Theorem 2.25 (Classical completeness for regular logic). *If \mathbb{T} is a regular theory, and σ is a regular sequent over the signature of \mathbb{T} which is satisfied in all \mathbb{T} -models in **Sets**, then σ is provable in \mathbb{T} .*

Given a sequent $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$, we can replace this with the (clearly provably equivalent) sequent $\phi(\mathbf{x}) \vdash \phi(\mathbf{x}) \wedge \psi(\mathbf{x})$. This allows us to specify axiomatizations of theories of the form $\sigma \vdash \tau$, where $\tau \vdash \sigma$ is a tautology.

If every axiom in \mathbb{T} can be deduced in \mathbb{T}' , we say \mathbb{T}' is a *quotient* of the theory \mathbb{T} . We say two theories \mathbb{T} and \mathbb{T}' are equivalent if they are both quotients of one another. Each equivalence class of theories contains precisely one deductively closed theory.

We order the deductively closed theories by inclusion. That is, for deductively closed theories \mathbb{T} , \mathbb{T}' , we set $\mathbb{T} \leq \mathbb{T}'$ if and only if every sequent in \mathbb{T} is also in \mathbb{T}' ; this is equivalent to saying that \mathbb{T}' is a quotient of \mathbb{T} . Deductively closed theories are closed under arbitrary intersections, and therefore they form a complete lattice.

Let R be a ring. Denote by $\text{mod-}R$ the category of finitely presented right R -modules. The category of additive functors $\text{mod-}R \rightarrow \mathbf{Ab}$ will be denoted $(\text{mod-}R, \mathbf{Ab})$. Any regular formula $\phi(\mathbf{x})$ in the language of R -modules defines a functor in this category - it can readily be shown that for a given module M , the collection $\phi(M)$ of tuples $\mathbf{m} \in M$ which satisfy ϕ is closed under addition and includes the tuple $\mathbf{0}$, and so is an abelian group. This category forms an additive analogue of a classifying topos (classifying toposes are described in the next section). Given any Grothendieck abelian category \mathcal{A} , any additive functor $M : R \rightarrow \mathcal{A}$ can be extended to a functor $\llbracket - \rrbracket_M : (\text{mod-}R, \mathbf{Ab}) \rightarrow \mathcal{A}$ preserving limits and colimits, such that M factorizes as:

$$R \xrightarrow{y} \text{mod-}R \xrightarrow{y} (\text{mod-}R, \mathbf{Ab}) \xrightarrow{\llbracket - \rrbracket_M} \mathcal{A}$$

This follows in two steps: from [23, 10.2.37], the category of finitely presentable objects in $(\text{mod-}R, \mathbf{Ab})$, denoted $\mathbf{fp}(\text{mod-}R, \mathbf{Ab})$, is the *free abelian category* generated by R , that is, for any abelian category \mathcal{A} , any additive functor $R \rightarrow \mathcal{A}$ factors through the embedding $y^2 : R \rightarrow \mathbf{fp}(\text{mod-}R, \mathbf{Ab})$. However, there is a duality between small abelian categories and locally coherent abelian categories, see [29, 2.2]. Under this duality, the abelian category $\mathbf{fp}(\text{mod-}R, \mathbf{Ab})$ corresponds to the full category $(\text{mod-}R, \mathbf{Ab})$, and the property described above follows from this.

We will prove a result associating the finite-type localizations of $(\text{mod-}R, \mathbf{Ab})$

with the regular quotients of the theory of modules over R . As noted in section 2.4, the finite-type localizations of a locally coherent abelian category \mathcal{C} correspond to the Serre subcategories of the category of finitely presented objects in \mathcal{C} . For the category $(\text{mod-}R, \mathbf{Ab})$, there is a rather nice characterization of the finitely presentable objects, due to Burke [9].

A *positive primitive pair* over R , denoted ϕ/ψ , is a pair of regular formulas $\phi(\mathbf{x})$ and $\psi(\mathbf{x})$ in the language of modules over R , such that the sequent $\psi(\mathbf{x}) \vdash \phi(\mathbf{x})$ is deducible in the theory of modules over R . Given a different selection of variables \mathbf{y} , we regard the regular formulas $\phi(\mathbf{y})$ and $\psi(\mathbf{y})$ as defining the same positive primitive pair. A positive primitive pair ϕ/ψ defines a functor $\text{mod-}R \rightarrow \mathbf{Ab}$, mapping each module M to the quotient group $\phi(M)/\psi(M)$.

The *category of pp-pairs over R* , denoted \mathbf{L}_R^{eq+} , is a category having the positive primitive pairs over R as its objects. A morphism ϕ/ψ (both with free variables \mathbf{x}) to χ/η (both with free variables \mathbf{y}) is given a regular formula $\rho(\mathbf{x}, \mathbf{y})$ such that the following three sequents can be proved in the theory of modules over R :

1. $\rho(\mathbf{x}, \mathbf{y}) \wedge \phi(\mathbf{x}) \vdash \chi(\mathbf{y})$
2. $\rho(\mathbf{x}, \mathbf{y}) \wedge \psi(\mathbf{x}) \vdash \eta(\mathbf{y})$
3. $\phi(\mathbf{x}) \vdash \exists \mathbf{y} \rho(\mathbf{x}, \mathbf{y})$

These conditions state precisely that for any R -module M , the definable set $\rho(M)$ will be the graph of a map between the quotient groups $\phi(M)/\psi(M)$ and $\chi(M)/\eta(M)$.

Lemma 2.26. ([23, 3.2.10]) *With the above definition, \mathbf{L}_R^{eq+} is an abelian category.*

Theorem 2.27. ([23, 10.2.30]) *For a ring R , the category \mathbf{L}_R^{eq+} is equivalent to the category of finitely presentable objects in the functor category $(\text{mod-}R, \mathbf{Ab})$.*

Now let S be a Serre subcategory of \mathbf{L}_R^{eq+} . We can associate to S a collection of sequents over the regular theory of modules over R . Each object in \mathbf{L}_R^{eq+} is of the form ϕ/ψ , where ϕ and ψ are pp-formulas over R , and for every module M , $\psi(M) \subseteq \phi(M)$. Syntactically, this is equivalent to saying that the sequent $\psi(\mathbf{x}) \vdash \phi(\mathbf{x})$ is a tautology

in the theory (by classical completeness). We associate to each object ϕ/ψ of \mathbf{L}_R^{eq+} the regular sequent $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$. To each Serre subcategory S of \mathbf{L}_R^{eq+} , we associate the collection $T(S)$ of all the sequents corresponding to objects in S .

Conversely, suppose we are given a quotient \mathbb{T} of the regular theory of modules over R , that is, a collection of sequents of the form $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$, where $\psi(\mathbf{x}) \vdash \phi(\mathbf{x})$ is a tautology in the theory. We can associate to this the subcategory $S(\mathbb{T})$ of \mathbf{L}_R^{eq+} consisting of all those pp-pairs ϕ/ψ , where $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$ is a sequent in \mathbb{T} .

Theorem 2.28. *Let R be a ring. For each Serre subcategory S of \mathbf{L}_R^{eq+} , $T(S)$ is closed under deductions. For each deductively closed quotient \mathbb{T} of the regular theory of modules over R , $S(\mathbb{T})$ is a Serre subcategory. This sets up a bijective correspondence between quotients of the regular logic and Serre subcategories of \mathbf{L}_R^{eq+} .*

Proof. Let S be a Serre subcategory of \mathbf{L}_R^{eq+} . We show that $T(S)$ is deductively closed. We verify this for each of the rules of regular logic.

Let $\phi(\mathbf{x})$ and $\psi(\mathbf{x})$ be any pp-formulas; it follows from Classical Completeness that the following pp-pairs are isomorphic to 0 in \mathbf{L}_R^{eq+} , and therefore are contained in S : ϕ/ϕ , $\top/(x = x)$, $(\mathbf{x} = \mathbf{y}) \wedge \phi(\mathbf{x})/\phi(\mathbf{y})$.

Furthermore, for the substitution rule, the formula $\mathbf{x} = \mathbf{y}$ provides a pp-definable map $\phi(\mathbf{x})/\psi(\mathbf{x}) \rightarrow \phi(\mathbf{y})/\psi(\mathbf{y})$.

To verify that the cut rule holds, suppose ϕ/ψ and ψ/χ are objects in S . There is an exact sequence

$$0 \rightarrow \psi/\chi \rightarrow \phi/\chi \rightarrow \phi/\psi \rightarrow 0$$

in \mathbf{L}_R^{eq+} ; it follows that ϕ/χ is in S .

To verify the rule for finite conjunction, suppose ϕ/ψ and ϕ/χ are in S . Consider the map $\phi \rightarrow \phi/\psi \oplus \phi/\chi$ given on each module M by $x \mapsto (x + \psi(M), x + \chi(M))$. The kernel of this map will always be $\psi \wedge \chi(M)$, so the kernel in \mathbf{L}_R^{eq+} is $\psi \wedge \chi$. Thus the image $\phi/\psi \wedge \chi$ is a subobject of the product $\phi/\psi \oplus \phi/\chi$, and so this image is in S .

Suppose $\phi(\mathbf{x}, y)$ and $\psi(\mathbf{x})$ are formulas. The pp-pairs $\phi(\mathbf{x}, y)/\psi(\mathbf{x}, y)$ and $(\exists y)\phi(\mathbf{x}, y)/\psi(\mathbf{x})$ are isomorphic in \mathbf{L}_R^{eq+} via the pp-definable morphism $\mathbf{x} = \mathbf{x}'$.

This proves the first half of the claim.

Let \mathbb{T} be a quotient of the regular theory of R -modules.

Let ϕ/ψ be an object in $S(\mathbb{T})$; we show that for any pp-pair θ/χ admitting a (pp-definable) isomorphism $\rho : \phi/\psi \rightarrow \chi/\theta$, there is a deduction of $\theta(\mathbf{y}) \vdash \chi(\mathbf{y})$ from $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$.

The pp-formula $\rho(\mathbf{x}, \mathbf{y})$ defines a surjection $\phi/\psi \rightarrow \chi/\theta$. By the completeness theorem for regular logic [17, D1.5.4] the following sequents, corresponding to the implications stated on p.92 of [23], can be proven in the regular theory of modules over R :

1. $\rho(\mathbf{x}, \mathbf{y}) \wedge \phi(\mathbf{x}) \vdash \chi(\mathbf{y})$
2. $\rho(\mathbf{x}, \mathbf{y}) \wedge \psi(\mathbf{x}) \vdash \theta(\mathbf{y})$
3. $\phi(\mathbf{x}) \vdash \exists \mathbf{y} \rho(\mathbf{x}, \mathbf{y})$

Furthermore, since ρ is stated to be a surjection, the sequent $\chi(\mathbf{y}) \vdash \exists \mathbf{x} \rho(\mathbf{x}, \mathbf{y}) \wedge \phi(\mathbf{x})$ is true in all R -modules, and is therefore provable in the theory.

We use the above sequents to write out the following deduction of $\chi(\mathbf{y}) \vdash \theta(\mathbf{y})$ from $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$.

$$\begin{array}{c}
 \frac{\frac{\frac{\rho(\mathbf{x}, \mathbf{y}) \wedge \phi(\mathbf{x}) \vdash \phi(\mathbf{x}) \quad \phi(\mathbf{x}) \vdash \psi(\mathbf{x})}{\rho(\mathbf{x}, \mathbf{y}) \wedge \phi(\mathbf{x}) \vdash \psi(\mathbf{x})}}{\rho(\mathbf{x}, \mathbf{y}) \wedge \phi(\mathbf{x}) \vdash \rho(\mathbf{x}, \mathbf{y})}}{\rho(\mathbf{x}, \mathbf{y}) \wedge \phi(\mathbf{x}) \vdash \rho(\mathbf{x}, \mathbf{y}) \wedge \psi(\mathbf{x})} \quad \rho(\mathbf{x}, \mathbf{y}) \wedge \psi(\mathbf{x}) \vdash \theta(\mathbf{y})}{\rho(\mathbf{x}, \mathbf{y}) \wedge \phi(\mathbf{x}) \vdash \theta(\mathbf{y})} \\
 \frac{\chi(\mathbf{y}) \vdash \exists \mathbf{x} \rho(\mathbf{x}, \mathbf{y}) \wedge \phi(\mathbf{x}) \quad \exists \mathbf{x} \rho(\mathbf{x}, \mathbf{y}) \wedge \phi(\mathbf{x}) \vdash \theta(\mathbf{y})}{\chi(\mathbf{y}) \vdash \theta(\mathbf{x})}
 \end{array}$$

The converse follows similarly, using the surjection $\rho^{-1} : \chi/\theta \rightarrow \phi/\psi$.

Now suppose we are given any exact sequence

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$$

in \mathbf{L}_R^{eq+} . We can represent the object B as some pp-pair ϕ/χ . Replacing the object A

with its image under f , we can express it as a pp-pair ψ/χ , with the graph of f being the diagonal in $\psi \times \psi$. The cokernel C is then given by the pp-pair ϕ/ψ . Now if we have a Serre subcategory containing A and C , then the corresponding theory contains the sequents $\psi(\mathbf{x}) \vdash \chi(\mathbf{x})$ and $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$. To obtain the sequent corresponding to B , $\phi(\mathbf{x}) \vdash \chi(\mathbf{x})$, one just needs a single application of the cut rule.

Furthermore, if B is in the Serre subcategory, we know that the sequent $\phi(\mathbf{x}) \vdash \chi(\mathbf{x})$ is in the corresponding theory. Since $\psi(\mathbf{x}) \vdash \phi(\mathbf{x})$ and $\chi(\mathbf{x}) \vdash \psi(\mathbf{x})$ can both be deduced in the theory, the sequents $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$ and $\psi(\mathbf{x}) \vdash \chi(\mathbf{x})$, corresponding to A and C respectively, can now be deduced, again using the cut rule. \square

2.6 Models of a geometric theory

In the previous section, we looked at the regular sequents in the language of a many-sorted signature of algebras Σ , and defined what their models were in **Sets**. In this section, we will introduce broader classes of formulas, and describe their models in arbitrary Grothendieck toposes.

Let Σ be a many-sorted signature. A regular formula over a many sorted signature is one of the form $\exists \mathbf{y} \wedge_{i=1}^n \psi_i(\mathbf{x}, \mathbf{y})$ where again, the formulas ψ_i are assumed to be atomic formulas. For a signature that includes relations, atomic formulas are either equations in the variables, or formulas of the form $R(z_1, \dots, z_n)$, where R is a relation symbol in Σ . A *coherent formula* over Σ is a formula of the form $\bigvee_{i=1}^n \phi_i(\mathbf{x})$ where each formula $\phi_i(\mathbf{x})$ is a regular formula. A *geometric formula* is a formula of the form $\bigvee_{i \in I} \phi_i(\mathbf{x})$, where I can be a set of arbitrary size. Note that allowing I to be empty enables us to add the formula \perp ('false') to the language.

As in the additive case, a geometric formula $\phi(\mathbf{x})$ over a signature Σ can be thought of as a functor $\phi : \mathbf{Str}(\Sigma) \rightarrow \mathbf{Sets}$, where each Σ -structure X is mapped to $\phi(X)$, the set of tuples \mathbf{x} from X which satisfy the formula $\phi(\mathbf{x})$. If $f : X \rightarrow X'$ is a map of Σ -structures and $\mathbf{x} \in \phi(X)$, then $f(\mathbf{x}) \in \phi(X')$, since $\phi(\mathbf{x})$ is a positive formula.

Coherent and geometric formulas admit their own logic, with the additional rules

of deduction for disjunction. The rules of deduction for coherent logic are the rules (a), (b), (c), (d) and (f), where (d) consists of the rule for finitary disjunction described below:

(d)

$$\perp \vdash \phi(\mathbf{x}), \quad \phi(\mathbf{x}) \vdash \phi(\mathbf{x}) \vee \psi(\mathbf{x}), \quad \psi(\mathbf{x}) \vdash \phi(\mathbf{x}) \vee \psi(\mathbf{x})$$

and the rule

$$\frac{\phi(\mathbf{x}) \vdash \chi(\mathbf{x}) \quad \psi(\mathbf{x}) \vdash \chi(\mathbf{x})}{\phi(\mathbf{x}) \vee \psi(\mathbf{x}) \vdash \chi(\mathbf{x})}.$$

For geometric logic, we add the rule (h) for infinitary disjunction:

(h)

$$\frac{(\phi_i(\mathbf{x}) \vdash \chi(\mathbf{x}))_{i \in I}}{\bigvee_{i \in I} \phi_i(\mathbf{x}) \vdash \chi(\mathbf{x})}.$$

Locally finitely presentable categories can be characterized as categories of models of *finite limit theories*, that is, theories whose axioms are all of the form

$$\phi(\mathbf{x}) \vdash (\exists! \mathbf{y})\psi(\mathbf{x}, \mathbf{y})$$

where ϕ and ψ are finite conjunctions of atomic formulas. It is clear why a finitary essentially algebraic category is of this form - for each partial operation σ , we can write a formula ϕ_σ encoding the definability condition for σ , and then $\sigma(\mathbf{x})$ will correspond to the element y whose existence we can assert by this formula.

These limit theories correspond to the cartesian theories presented in [17, D1.3.4], though the presentation there is slightly different. Johnstone defines a cartesian formula relative to a (regular) theory \mathbb{T} as follows: atomic formulas are cartesian, cartesian formulas are closed under finite composition and whenever the sequent $\phi(\mathbf{x}, y) \wedge \phi(\mathbf{x}, z) \vdash (y = z)$ is provable in \mathbb{T} , then the formula $\exists! y \phi(\mathbf{x}, y)$ is cartesian. It is clear that a cartesian theory can be written as a finite limit theory if we prefer; to write a finite limit theory as a cartesian theory, we need to write the theory out without the unique existence quantifier, and add axioms of the form $\phi(\mathbf{x}) \wedge \psi(\mathbf{x}, y) \wedge$

$\psi(\mathbf{x}, z) \vdash (y = z)$.

Models of a geometric theory \mathbb{T} can be defined in any Grothendieck topos \mathcal{E} , not just in the topos **Sets**. Let \mathbb{T} be a geometric theory over an underlying S -sorted signature Σ . A Σ -structure M in \mathcal{E} is specified by the following:

1. for each sort X of S , an object MX of \mathcal{E} .
2. for each function symbol $f : X_1 \times \dots \times X_n \rightarrow X$, a morphism $Mf : MX_1 \times \dots \times MX_n \rightarrow MX$ in \mathcal{E} .
3. for each relation symbol $R \rightarrow X_1 \times \dots \times X_n$ in Σ , a subobject $MR \rightarrow MX_1 \times \dots \times X_n$ in \mathcal{E} .

If M and N are Σ -structures, a Σ -homomorphism $h : M \rightarrow N$ is specified by a collection of morphisms $h_X : MX \rightarrow NX$ in \mathcal{E} , such that

4. for each function symbol $f : X_1 \times \dots \times X_n \rightarrow X$, the diagram below commutes:

$$\begin{array}{ccc} MX_1 \times \dots \times MX_n & \xrightarrow{Mf} & MX \\ \downarrow h_{X_1} \times \dots \times h_{X_n} & & \downarrow h_X \\ NX_1 \times \dots \times NX_n & \xrightarrow{Nf} & NX \end{array}$$

5. for each relation symbol $R \rightarrow X_1 \times \dots \times X_n$, there is a commutative diagram of the form

$$\begin{array}{ccc} MR & \longrightarrow & MX_1 \times \dots \times MX_n \\ \downarrow & & \downarrow h_{X_1} \times \dots \times h_{X_n} \\ NR & \longrightarrow & NX_1 \times \dots \times NX_n \end{array}$$

In any Σ -structure M in any Grothendieck topos \mathcal{E} , we can define *interpretations* of terms and formulas. A term $t(x_1, \dots, x_n)$ of sort X in variables $x_1 \in X_1, \dots, x_n \in X_n$ is interpreted by a morphism

$$\llbracket \mathbf{x} : t \rrbracket_M : MX_1 \times \dots \times MX_n \rightarrow MX.$$

- if the term is x_i , for some $i = 1, \dots, n$, then the interpretation of the term is the i th projection map $MX_1 \times \dots \times MX_n \rightarrow MX_i$.

- if the term is $f(t_1, \dots, t_m)$, where each term t_i is a variable in (x_1, \dots, x_n) of sort Y_i , then the interpretation $\llbracket \mathbf{x} : t \rrbracket_M$ of the term is the composite

$$MX_1 \times \dots \times MX_n \xrightarrow{\llbracket \mathbf{x}:t_1 \rrbracket_M \times \dots \times \llbracket \mathbf{x}:t_m \rrbracket_M} MY_1 \times \dots \times MY_m \xrightarrow{f} MX.$$

A geometric formula $\phi(x_1, \dots, x_n)$ in variables $x_1 \in X_1, \dots, x_n \in X_n$ is interpreted by a subobject

$$\llbracket \mathbf{x} : \phi \rrbracket_M \hookrightarrow MX_1 \times \dots \times MX_n.$$

- if ϕ is an atomic formula of the form $t_1(\mathbf{x}) = t_2(\mathbf{x})$ for terms t_1 and t_2 of sort X , then the interpretation of ϕ is the equaliser in \mathcal{E} of the maps $\llbracket \mathbf{x} : t_1 \rrbracket_M, \llbracket \mathbf{x} : t_2 \rrbracket_M : MX_1 \times \dots \times MX_n \rightarrow MX$.
- if ϕ is an atomic formula of the form $R(t_1, \dots, t_m)$ where each t_i is a term of sort Y_i , then the interpretation of ϕ is given by the pullback

$$\begin{array}{ccc} \llbracket \mathbf{x} : \phi(\mathbf{t}) \rrbracket_M & \xrightarrow{\quad\quad\quad} & MR \\ \downarrow & & \downarrow \\ MX_1 \times \dots \times MX_n & \xrightarrow{\llbracket \mathbf{x}:t_1 \rrbracket_M \times \dots \times \llbracket \mathbf{x}:t_m \rrbracket_M} & MY_1 \times \dots \times MY_m \end{array}$$

- if ϕ is the formula \perp , it is interpreted by the map $0 \rightarrow MX_1 \times \dots \times MX_n$; if it is the formula \top , it is interpreted by the identity map $1_{MX_1 \times \dots \times MX_n} : MX_1 \times \dots \times MX_n \rightarrow MX_1 \times \dots \times MX_n$.
- the logical connectives \wedge and \vee are just interpreted as the corresponding operators in the subobject lattice $\text{Sub}(MX_1 \times \dots \times MX_n)$, which in a Grothendieck topos is a complete Heyting algebra.
- if ϕ is the formula $\exists y \psi(\mathbf{x}, y)$ for some variable y of sort Y , then the interpretation of ϕ is the image factorization of the map

$$\llbracket (\mathbf{x}, y) : \psi \rrbracket_M \longrightarrow MX_1 \times \dots \times MX_n \times MY \xrightarrow{\pi} MX_1 \times \dots \times MX_n$$

where π is the projection map onto the first n coordinates.

Now let M be a Σ -structure in a Grothendieck topos \mathcal{E} . The sequent $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$ is said to *hold* in M if $\llbracket \mathbf{x} : \phi \rrbracket_M \leq \llbracket \mathbf{x} : \psi \rrbracket_M$ in the subobject lattice $\text{Sub}(MX_1 \times \dots \times MX_n)$.

Let \mathcal{T} be a geometric theory. A Σ -structure M in a Grothendieck topos \mathcal{E} is said to be a *model* of \mathbb{T} if every sequent in \mathbb{T} holds in M .

It can easily be checked that the deductions for geometric logic are valid: that is, for each of the rules (a), (b), (c), (d), (f), (h), (i), if the premises for that rule hold in a Σ -structure M , then the conclusion holds as well.

If the signature Σ is algebraic (i.e., contains no relation symbols), and all the sequents in the theory \mathbb{T} are of the form $\top \vdash \phi(\mathbf{x})$ where $\phi(\mathbf{x})$ is an equation between terms in the variables \mathbf{x} , then the theory \mathbb{T} is precisely an equational theory (Σ, E) , as described in section 2.1. In this case, we say that the models of the theory \mathbb{T} are (Σ, E) -*objects* in \mathcal{E} . This terminology is commonly used for widely studied equational theories such as groups or rings. If for example, (Σ, E) is the equational theory of groups, we refer to its models in a topos \mathcal{E} as *group objects* in \mathcal{E} (and similarly for rings, monoids, etc.).

2.7 Classifying toposes

In this section, we associate with any geometric theory \mathbb{T} a *classifying topos* $\mathbf{Set}[\mathbb{T}]$, which has the property that the models of \mathbb{T} in any topos \mathcal{E} correspond precisely with the geometric morphisms $\mathcal{E} \rightarrow \mathbf{Set}[\mathbb{T}]$. This works because the objects in $\mathbf{Set}[\mathbb{T}]$ can be thought of as (disjoint unions of) formulas in the theory, and so given a topos \mathcal{E} and any model M of \mathbb{T} in \mathcal{E} , we can send each formula in the classifying topos to the interpretation of that formula in M , giving us the inverse image part of the geometric morphism.

There are a number of ways of constructing the classifying topos; these are detailed in [17, D.3.1]. The approach we use will first construct the classifying topos for a cartesian theory; if a geometric theory \mathbb{T} is a quotient of a geometric theory \mathbb{T}_0 , then $\mathbf{Set}[\mathbb{T}]$ can be constructed as a localization of $\mathbf{Set}[\mathbb{T}_0]$. In particular, any geometric

theory \mathbb{T} over a signature Σ is a quotient of the (cartesian) empty theory over Σ , so this is enough to construct the classifying topos for all geometric theories \mathbb{T} .

Let \mathbb{T} be a cartesian theory. We construct a category over \mathbb{T} , called the *syntactic site*, $\mathcal{C}_{\mathbb{T}}$. The objects of $\mathcal{C}_{\mathbb{T}}$ are equivalence classes of cartesian formulas over the signature of \mathbb{T} , where two formulas ϕ and ψ are taken to be equivalent if there is a derivation of ψ from $\mathbb{T} \cup \{\phi\}$ and vice versa. Given a cartesian formula ϕ over the signature, we write $[\phi]$ for the equivalence class of formulas containing ϕ .

The arrows in $\mathcal{C}_{\mathbb{T}}$ are (cartesian-)definable maps between the formulas - that is, the arrows $[\phi] \rightarrow [\psi]$ are equivalence classes of cartesian formulas θ such that for any model M of \mathbb{T} , θ defines a map $\phi(X) \rightarrow \psi(X)$, with two such formulas θ and θ' being equivalent if they define the same map in every model.

Theorem 2.29. *[17, D1.4.7] Let \mathbb{T} be a Cartesian theory. Models of the theory in \mathcal{E} are equivalent to functors $\mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{E}$ preserving finite limits.*

We will use $\mathcal{C}_{\mathbb{T}}$ as the base category for the underlying site of the classifying topos. We will need the following result, from [20, VII.9.4]. For a site (\mathcal{C}, J) and a cocomplete category \mathcal{E} , call a functor $F : \mathcal{C} \rightarrow \mathcal{E}$ *J-continuous* if every covering sieve in \mathcal{C} is mapped to an epimorphic family in \mathcal{E} .

Theorem 2.30. *[20, VII.9.4] Let (\mathcal{C}, J) be a site where \mathcal{C} has finite limits, and let \mathcal{E} be a Grothendieck topos. There is an equivalence of categories between geometric morphisms $\mathcal{E} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$ and J-continuous functors $\mathcal{C} \rightarrow \mathcal{E}$ preserving finite limits.*

We give a brief outline of why this is true, and refer the reader to [20] for details. For the small category \mathcal{C} , the Yoneda embedding $y : \mathcal{C} \rightarrow \mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ is a free cocompletion of \mathcal{C} ; that is, every object in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ is a colimit of representable objects. Thus, given a functor $F : \mathcal{C} \rightarrow \mathcal{E}$ with \mathcal{E} cocomplete, we can extend this to $\tilde{F} : \mathbf{Sets}^{\mathcal{C}^{\text{op}}} \rightarrow \mathcal{E}$ by simply stipulating that \tilde{F} preserve colimits: an object P in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ can be presented as the colimit of a diagram $D : \mathcal{D} \rightarrow \mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ whose objects are representable functors, i.e., D factors through the Yoneda embedding via some functor $D' : \mathcal{D} \rightarrow \mathcal{C}$. We define $\tilde{F}(P)$ to be the colimit in \mathcal{E} of the diagram $F.D' : \mathcal{D} \rightarrow \mathcal{E}$. It can be shown

that this extension of F is well defined; this construction is the *left Kan extension* of F along y , cf [7, 3.7.2]. The extension inherits the property of preserving finite limits from F . It maps the covering sieves in J to epimorphic families in \mathcal{E} ; this translates into mapping the corresponding inclusions in $\mathbf{Sets}^{\text{cop}}$ to isomorphisms in \mathcal{E} , so the functor \tilde{F} factors through the localization $a_J : \mathbf{Sets}^{\text{cop}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$. The functor $F^* : \mathbf{Sh}(\mathcal{C}, J) \rightarrow \mathcal{E}$ so constructed has a right adjoint F_* , by the special adjoint functor theorem.

We see from the above two theorems how to construct the classifying topos for a cartesian theory \mathbb{T} . This construction is detailed in [17, D3.1]. For a given topos \mathcal{E} , \mathbb{T} -models in \mathcal{E} correspond to functors $F : \mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{E}$ preserving finite limits, which in turn correspond to geometric morphisms $(F^*, F_*) : \mathcal{E} \rightarrow (\mathcal{C}_{\mathbb{T}}, \mathbf{Sets})$. The classifying topos of \mathbb{T} is given by the functor category $(\mathcal{C}_{\mathbb{T}}, \mathbf{Sets})$.

If \mathbb{T} is an arbitrary geometric theory, set \mathbb{T}_0 to be the empty theory over the same signature. Then every axiom in \mathbb{T} can be written in the form

$\phi(\mathbf{x}) \vdash \bigvee_{i \in I} (\exists \mathbf{y}) \psi_i(\mathbf{x}, \mathbf{y})$. Each such axiom corresponds to a family of morphisms in $\mathcal{C}_{\mathbb{T}_0}$:

$$[\psi_i] \xrightarrow{\psi_i \wedge \mathbf{x} = \mathbf{x}'} \phi[\mathbf{x}' / \mathbf{x}].$$

Moreover, the sequent holds in a model M of the theory precisely when this family of morphisms is mapped to an epimorphic family by the interpretation functor $\llbracket - \rrbracket_M$. We denote by $J_{\mathbb{T}}$ the Grothendieck topology on $\mathcal{C}_{\mathbb{T}_0}$ generated by these families of morphisms (that is, for each axiom $\phi(\mathbf{x}) \vdash \bigvee_{i \in I} (\exists \mathbf{y}) \psi_i(\mathbf{x}, \mathbf{y})$, we take the associated family of morphisms to be covering).

Now given a topos \mathcal{E} , a model of \mathbb{T} is a model of \mathbb{T}_0 in which all the axioms in \mathbb{T} hold; this is equivalent to demanding that the families of mappings corresponding to each axiom $\phi(\mathbf{x}) \vdash \bigvee_{i \in I} (\exists \mathbf{y}) \psi_i(\mathbf{x}, \mathbf{y})$ are jointly epimorphic. So again by Theorem 2.30, the models of the geometric theory \mathbb{T} correspond to geometric morphisms $\mathcal{E} \rightarrow \mathbf{Sh}(\mathcal{C}_{\mathbb{T}_0}, J_{\mathbb{T}})$. Thus the classifying topos for \mathbb{T} is given by $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}_0}, J_{\mathbb{T}})$.

It can be shown that closing a collection of axioms with respect to \mathbb{T} is the same as taking the Grothendieck topology generated by the corresponding families of maps,

leading to the following result.

Theorem 2.31. [10, 3.6] *Let \mathbb{T} be a geometric theory over a signature Σ . Then associating a collection of axioms with the associated Grothendieck topology on \mathbb{T}_0 defines a 1-1 correspondence between the geometric quotients of \mathbb{T} and the localizations of the classifying topos of \mathbb{T} .*

Remark. We note in particular, from [11, 5.5.8], that if \mathbb{T} is regular theory, this correspondence restricts to a correspondence between the regular quotients of \mathbb{T} and the Grothendieck topologies J on $\mathcal{C}_{\mathbb{T}}$ such that every covering sieve contains a covering sieve generated by a single arrow.

Chapter 3

Sheaves as essentially algebraic objects

3.1 Sheaves as essentially algebraic objects

Let \mathcal{C} be a small category. The category $(\mathcal{C}, \mathbf{Sets})$ of set valued functors on \mathcal{C} is described by a multi-sorted equational theory, which we will denote $\Gamma_{\mathcal{C}}$. This theory is described as follows:

- For each object C of \mathcal{C} , we take a corresponding sort \mathbf{C} .
- For each morphism $f : C \rightarrow C'$ in \mathcal{C} , take a function symbol $\mathbf{f} : \mathbf{C} \rightarrow \mathbf{C}'$.
- For each commutative diagram in \mathcal{C} of the form

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ & \searrow h & \downarrow g \\ & & C \end{array}$$

add an equation to E in one variable \mathbf{x} of sort \mathbf{A} stating $\mathbf{g}\mathbf{f}(\mathbf{x}) = \mathbf{h}(\mathbf{x})$.

If (\mathcal{C}, J) is a site, we can describe the presheaves on \mathcal{C} as the set-valued functors on \mathcal{C}^{op} in the manner just described. Furthermore, we can extend the theory to an essentially algebraic theory $(\Sigma, E, \Sigma_t, \text{Def})$ whose models are the sheaves on the site.

- For every covering sieve J of an object \mathcal{C} , we take a partial operation

$$\sigma_J : \prod_{f \in J} \text{dom}(f) \rightarrow \mathcal{C}$$

, in variables $\mathbf{x} = (\mathbf{x}_f)_{f \in J}$ (each variable (\mathbf{x}_f) is of sort $\text{cod}(f)$).

- The equations in $\text{Def}(\sigma_J)$ are those of the form $\mathbf{g}f(\mathbf{x}_f) = \mathbf{h}(\mathbf{x}_h)$ whenever we have a commutative diagram

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ & \searrow h & \downarrow g \\ & & C \end{array}$$

such that g and h are in J .

- We add equations to E stating that for each $f \sigma_J(\mathbf{x}) = \mathbf{x}_f$ for every covering sieve J and every $f \in J$.

It is easily checked that the models of this essentially algebraic theory are just the sheaves for the topology. We write $\Gamma_{(\mathcal{C}, J)}$ for this essentially algebraic theory.

In fact it suffices to take a basis for the Grothendieck topology in the above. That is, given a basis K for J , define an essentially algebraic theory by taking the partial operations $\sigma_{\{f_i\}}$ defined for each covering family $\{f_i\}$ of morphisms in K as above. The models of this algebraic theory are again the sheaves for the topology. We denote this essentially algebraic theory by $\Gamma_{(\mathcal{C}, K)}$.

Let J be a topology, such that there exists a regular cardinal λ and a basis K for the topology such that every covering family in K has less than λ elements. Then every function symbol in $\Gamma_{(\mathcal{C}, K)}$ will take fewer than λ arguments, and by Lemma 2.15, the category $\mathbf{Sh}(\mathcal{C}, J)$ will be locally λ -generated.

3.2 Locally coherent and finitely presented toposes

By considering sheaves as essentially algebraic objects, we can understand the notions of finite presentability and coherence in a very concrete way. In this section, we will

use this to characterize the different local generation conditions for toposes. The results here are mostly known, see e.g., [4, VI.2], but this approach gives us a different way of thinking about them.

To provide characterizations of toposes with these various local generation properties, we start by introducing the following form of the Comparison Lemma.

Lemma 3.1. (*[20, p.589]*) *If \mathcal{C} is a full subcategory of the Grothendieck topos, \mathcal{E} whose objects form a generating set, and J is the topology on \mathcal{C} in which the covering sieves on an object C are precisely those containing an epimorphic family of morphisms, then \mathcal{E} is equivalent to $\mathbf{Sh}(\mathcal{C}, J)$.*

We use this result to find sites for a given topos. For instance, if we assume the topos \mathcal{E} is locally finitely generated, we can take \mathcal{C} to be the collection of finitely generated objects in \mathcal{E} . However, if J is a sieve on a finitely generated object in \mathcal{E} which contains an epimorphic family of morphisms, then in particular J contains a finite epimorphic family of morphisms. This observation leads to one half of the following result.

Proposition 3.2. *A topos \mathcal{E} is locally finitely generated if and only if it is equivalent to $\mathbf{Sh}(\mathcal{C}, J)$, for some site (\mathcal{C}, J) , where every sieve in the topology J contains a dense finitely generated sieve.*

Proof. It remains to show that a topos of this form is locally finitely generated. By Theorem 2.7, it suffices to prove that sheaves are closed under directed unions.

Let $\{d_{ij} : D_i \rightarrow D_j \mid i \leq j \in I\}$ be a directed system in $\mathbf{Sh}(\mathcal{C}, J)$, where (I, \leq) is a directed poset and each d_{ij} is a monomorphism. Then if $\{d_i : D_i \rightarrow D\}$ is a colimit cocone for the system, the object D must be a sheaf for J . For suppose we have an inclusion $s : S \rightarrow \text{Hom}(-, C)$ where S is a dense sieve on C , and a map $f : S \rightarrow D$. The sieve S contains a finitely generated dense sieve S' , with inclusion map $s' : S' \rightarrow S$, say. This gives us a map $f' = fs' : S' \rightarrow D$. Since S' is finitely generated, the map f' factors through $d_i : D_i \rightarrow D$ for some $i \in I$. But since D_i is a sheaf, we have a unique extension of f' to a map $\tilde{f} : \text{Hom}(-, C) \rightarrow D_i$.

We claim that $\tilde{f}s$ is a factorization of f through the colimit; that is, the $d_i\tilde{f}s = f$. If not, there is a map $x : X \rightarrow S$ with $d_i\tilde{f}s x \neq fx$. The image of X is a strong quotient of X , and so is finitely generated. The union of S' and the $\text{im}(X)$, denoted $S' \vee \text{im}(X)$, is the image of the map $S' \amalg \text{im}(X) \rightarrow S$ determined by s' and x , and so is also finitely generated. Denote the inclusion map $S' \vee \text{im}(X) \rightarrow S$ by m . The map fm factors through some map d_j in the colimit cocone, say as $fm = d_jg$. Since $S' \vee \text{im}(X)$ is dense, the map g can be extended to the whole of $\text{Hom}(-, C)$. Now consider the maps $d_{i, i \vee j}\tilde{f}$, $d_{j, i \vee j}g'$. These maps must be equal, since both are extensions of the map $d_{i, i \vee j}f'$ to $\text{Hom}(-, C)$. But this contradicts the assumption that $d_i\tilde{f}s x \neq fx$. \square

Definition 3.3. *If (\mathcal{C}, J) is a site where every sieve in the topology J contains a dense finitely generated sieve, we say the topology J is of finite type. More generally, if every sieve in the topology J contains a dense λ -generated sieve, we say the topology J is of λ -type.*

A reflection functor $r : \mathcal{C} \rightarrow \mathcal{A}$ is said to be of finite type (respectively, of λ -type) if the inclusion functor $i : \mathcal{A} \rightarrow \mathcal{C}$ preserves directed colimits (respectively, λ -directed unions).

There is some confusion here over whether, for a finite-type localization, the inclusion functor $i : \mathcal{A} \rightarrow \mathcal{C}$ is required to preserve directed unions, or all directed colimits. This confusion is caused partly by the fact that if the small category \mathcal{C} has pullbacks, the two definitions are equivalent. However, the definition of finite type topology given here is fairly universal, and in general, it is only equivalent to demanding that the inclusion functor preserve λ -directed unions. We shall call a localization such that the inclusion functor preserves all directed colimits a *coherent type* localization.

It is shown in e.g., [27, 3.15] that for a Grothendieck topology J on a small category \mathcal{C} with pullbacks, the topology J is of finite type if and only if the localization functor $a : \mathbf{Sets}^{\mathcal{C}^{\text{op}}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$ is.

The next result characterizes locally finitely presented toposes. To make the proof easier to follow, we make the following definition. Let \mathcal{C} be a topos, and let $f : A \rightarrow C$, $g : B \rightarrow C$ be a pair of arrows in \mathcal{C} with common codomain. A *square* in \mathcal{C} over f

and g is a commutative diagram of the form

$$\begin{array}{ccc} X & \xrightarrow{x_1} & B \\ x_2 \downarrow & & \downarrow g \\ A & \xrightarrow{f} & C \end{array} .$$

We denote this square by (X, x_1, x_2) . Given two squares (X, x_1, x_2) , (Y, y_1, y_2) , we say a factorisation of X through Y is a map $x' : X \rightarrow Y$ with the property that $y_1 x' = x_1$, $y_2 x' = x_2$.

Theorem 3.4. *A topos \mathcal{E} is locally finitely presented if and only if it is equivalent to $\mathbf{Sh}(\mathcal{C}, J)$, for some site (\mathcal{C}, J) , where every sieve in the topology contains a dense sieve S with the property that (1) S is generated by a finite collection of arrows S' ; and (2) every pair of arrows $f, g \in S'$ admits a finite collection of squares, $X_i = (X_i, x_1^i, x_2^i)$ with the property that every other square $Y = (Y, y_1, y_2)$ factors through one of the X_i .*

Proof. Suppose the site (\mathcal{C}, J) is of the form described. Take as a basis for the topology the finite collections of arrows generating each sieve S' as described in the statement of the theorem. Associate with this basis the essentially algebraic theory as described in section 3.1. For each sieve S' , the equations stating that the squares X_i commute are sufficient to describe the essentially algebraic theory. This theory is finitary, and the category of its models is therefore locally finitely presented. This category is of course just $\mathbf{Sh}(\mathcal{C}, J)$.

It remains to show that if \mathcal{E} is locally finitely presented, then it is equivalent to the category of sheaves on a site of this form. By the previous result, if \mathcal{E} is locally finitely presented, it is equivalent to the category of sheaves on the site (\mathcal{C}, J) , where \mathcal{C} is the category of finitely generated objects in \mathcal{E} , and J is the topology on \mathcal{C} generated by the families of morphisms that are epimorphic families in \mathcal{E} . We claim that if J is not of the above form, then the objects in \mathcal{C} are not finitely presented in \mathcal{E} .

Suppose there is some object C in \mathcal{C} and some finitely generated J -dense sieve S on C , which is generated by a finite family of morphisms $\{s_i : S_i \rightarrow C\}_{1 \leq i \leq n}$.

Suppose there is some pair s_i, s_j , which does not admit a finite family of squares with the property (2). Consider the collection of all squares (Y_k, y_1^k, y_2^k) over s_i, s_j .

Build up a directed system of functors as follows: each functor F is generated by a pair of elements $x_i \in S_i$, for each $1 \leq i \leq n$. The directed system consists of all functors with this generating set and containing finitely many of the relations satisfied by the generating set of elements s_i in S . Thus these functors form a directed system (the join of two functors F_1, F_2 in the system is the functor whose set of relations is just the union of those for F_1 and F_2). The colimit of this directed system is clearly S . However, there is no map from S to any of the functors F_i , since this would force F_i to satisfy extra relations. Since S is isomorphic to $\text{Hom}(-, C)$ in the sheaf category, this contradicts the assertion that C is finitely presented in \mathcal{E} . \square

The most obvious examples of sites which fulfil the condition given in Proposition 3.4 are those where the topology is trivial, i.e., for each object C in \mathcal{C} , $JC = \{\text{Hom}(-, C)\}$ (this is the wholly obvious fact that presheaf categories are locally finitely presented) and those where the category \mathcal{C} has pullbacks (this is the equally obvious fact that locally coherent toposes are locally finitely presented).

If (\mathcal{C}, J) is a site as described above, the finitely presented objects in $\mathbf{Sh}(\mathcal{C}, J)$ are described by the following result, which appears as [27, 3.16]. This can be seen as an immediate consequence of Lemma 2.12 applied to the description of sheaves as essentially algebraic objects given in section 3.1.

Theorem 3.5. *Let J be a λ -type Grothendieck topology on a category \mathcal{C} .*

a If F is a λ -generated sheaf, there is a λ -generated presheaf P such that $F \cong aP$.

b If F is a λ -presented sheaf, there is a λ -presented presheaf P such that $F \cong aP$.

To look at coherent and locally coherent toposes, we will need the following result.

Lemma 3.6. *Let \mathcal{C} be any category. Then the full subcategory of \mathcal{C} consisting of coherent objects in \mathcal{C} is closed under pullbacks.*

Proof. Suppose we are given a pullback diagram of the form

$$\begin{array}{ccc} A & \xrightarrow{\pi_1} & B \\ \pi_2 \downarrow & & \downarrow g \\ C & \xrightarrow{f} & D \end{array}$$

in which the objects B , C and D are all assumed to be coherent. Since D is coherent and B and C are finitely generated, we have that A is finitely generated also.

Now suppose we have a further pullback diagram

$$\begin{array}{ccc} X & \xrightarrow{p_1} & Y \\ p_2 \downarrow & & \downarrow h \\ Z & \xrightarrow{f} & A \end{array}$$

in which Y and Z are finitely generated. Then the diagram

$$\begin{array}{ccc} X & \xrightarrow{p_1} & Y \\ p_2 \downarrow & & \downarrow g\pi_1 h \\ Z & \xrightarrow{f\pi_2 k} & D \end{array}$$

is a pullback diagram also, and so X is finitely generated by coherence of D .

Thus A is coherent also. □

The next result characterizes locally coherent toposes; it will require quite a bit of work to prove. This characterization was originally shown in [4, VI.2.1]. The proof we give here uses the idea of a presheaf as a model of an algebraic theory.

Proposition 3.7. *A topos \mathcal{E} is locally coherent if and only if it is equivalent to $\mathbf{Sh}(\mathcal{C}, J)$, for some site (\mathcal{C}, J) , where \mathcal{C} is closed under pullbacks and every sieve in the topology J is generated by a finite number of arrows.*

Such a topos is always cocomplete, so it suffices to prove that the functors $a\mathrm{Hom}(-, C)$ are coherent in such a topos (recall that a is the associated sheaf functor defined on page 35). We will need to look at the notion of separated presheaves. A presheaf P on a site (\mathcal{C}, J) is *separated* if for any object C in \mathcal{C} , and any cover S of

\mathcal{C} , if $x, y \in PC$ such that for all $f : D \rightarrow C$ in S , we have that if $Pf(x) = Pf(y)$, then $x = y$. That is, P is separated if elements of P agree on a cover only if they are the same. A presheaf P is separated if and only if P^+ is a sheaf.

The separated presheaves form a reflective subcategory $\mathbf{Sep}(\mathcal{C})$ of $\mathbf{Sets}^{\mathcal{C}^{op}}$, and the associated sheaf functor factors through this reflection. On any presheaf P , we define, for each object C of \mathcal{C} an equivalence relation R on PC given by

$$xRy \text{ if and only if } \exists S \in \mathcal{J}C \text{ such that } \forall f \in S, Pf(x) = Pf(y).$$

Given a map $f : C \rightarrow C'$, the corresponding map $Pf : PC' \rightarrow PC$ respects this equivalence relation, so this defines a functor $(-)_{\text{sep}} : \mathbf{Sets}^{\mathcal{C}^{op}} \rightarrow \mathbf{Sep}(\mathcal{C})$. We write P_{sep} for the image of P under this functor. This functor is a localization, so in particular the associated sheaf functor can be represented as $(-)_{\text{sep}}$ followed by one application of the plus-functor from section 2.3. The details of this can be found in, for example, [27, p.32].

If F is a presheaf on a site (\mathcal{C}, J) , we say a subpresheaf $s : S \rightarrow F$ is *dense* if as is an isomorphism. The dense subpresheaves of a presheaf P are closed under intersections, so they form a directed system, denoted $\mathbf{D}(P)$. If F and G are presheaves then a map between $f : aF \rightarrow aG$ may be represented by a map $f' : F' \rightarrow G_{\text{sep}}$ such that $af' = f$. Two maps $f : F' \rightarrow G_{\text{sep}}$ and $g : F'' \rightarrow G_{\text{sep}}$ represent the same map $aF' \rightarrow aG$ if they agree on some dense subobject of $F' \cap F''$.

Lemma 3.8. ([27, 3.9]) *For presheaves F and G on a site (\mathcal{C}, J) , there is a natural isomorphism*

$$\text{Hom}_{\mathbf{Sh}(\mathcal{C}, J)}(aF, aG) \cong \lim_{\rightarrow P' \in \mathbf{D}(P)} \text{Hom}_{\mathbf{Sets}^{\mathcal{C}^{op}}}(P', Q_{\text{sep}}).$$

Lemma 3.9. *Let (\mathcal{C}, J) be a site where \mathcal{C} is closed under pullbacks and every sieve in the topology J is generated by a finite number of arrows. Let $a : \mathbf{Sets}^{\mathcal{C}^{op}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$ be the localization functor. Then the functors $a\text{Hom}(-, C)$ are coherent objects in the sheaf category.*

Proof. The functor $a\text{Hom}(-, C)$ is finitely generated in the sheaf category by Theorem 3.5. It remains to prove the pullback property. We will show this by looking at the separated presheaf $\text{Hom}(-, C)_{\text{sep}}$. For each object C' in \mathcal{C} , the elements of $\text{Hom}(-, C)_{\text{sep}}(C')$ are equivalence classes of maps $f : C' \rightarrow C$, where f is equivalent to f' if there is some cover $\{g_i : G_i \rightarrow C'\}_{i \in I}$ such that $fg_i = f'g_i$, for all i .

Suppose we are given maps $\alpha : A \rightarrow a\text{Hom}(-, C)$, $\beta : B \rightarrow yC$ in the sheaf category, with A and B finitely generated. Since A and B are finitely generated in the sheaf category, they are isomorphic to sheaves aA' , aB' for some finitely generated presheaves A' and B' , by Theorem 3.5. The maps α and β can be represented by maps $\tilde{\alpha} : A^* \rightarrow \text{Hom}(-, C)_{\text{sep}}$ and $\tilde{\beta} : B^* \rightarrow \text{Hom}(-, C)_{\text{sep}}$ in $\mathbf{Sets}^{C^{\text{op}}}$ with A^* and B^* finitely generated dense subobjects of A' and B' respectively (we may assume A^* and B^* are finitely generated because the localization is of finite type). The presheaves A^*_{sep} and B^*_{sep} are finitely generated objects in the category $\mathbf{Sets}^{C^{\text{op}}}$, since they are quotients of the finitely generated objects A^* and B^* respectively.

Thus the map α and β are given by maps $\alpha^* : A^*_{\text{sep}} \rightarrow \text{Hom}(-, C)_{\text{sep}}$, $\beta^* : B^*_{\text{sep}} \rightarrow \text{Hom}(-, C)_{\text{sep}}$, where A^*_{sep} and B^*_{sep} are finitely generated objects in $\mathbf{Sets}^{C^{\text{op}}}$, and $a(\alpha^*) = \alpha$, $a(\beta^*) = \beta$.

Since A^*_{sep} and B^*_{sep} are objects in the presheaf category, we can assume they are generated as models of the algebraic theory $\Gamma_{\mathcal{C}}$ by elements $a_i \in A^*_{\text{sep}}C_i$ and $b_j \in B^*_{\text{sep}}C_j$. The map $\alpha^* : A^*_{\text{sep}} \rightarrow \text{Hom}(-, C)_{\text{sep}}$ identifies each of the generators a_i with an equivalence class of arrows $C_i \rightarrow C$, and we choose a representative $a_i^* : C_i \rightarrow C$ for each equivalence class. Similarly, we choose representatives b_j^* for the image of each of the generators b_j .

For each pair of generators a_i of A^*_{sep} and b_j of B^*_{sep} , take the pullback square

$$\begin{array}{ccc} C_i \times_C C_j & \xrightarrow{\pi_{i,j}^2} & C_j \\ \pi_{i,j}^1 \downarrow & & \downarrow b_j^* \\ C_i & \xrightarrow{a_i^*} & C \end{array}$$

Now define P to be the presheaf defined by taking generators (p_i, p_j) of sort $C_i \times_C C_j$ for each pair of generators a_i, b_j . Every term $t(p_i, p_j)$ of sort C' corresponds

to a map $t : C' \rightarrow C_i \times_C C_j$, which then corresponds to a pair of maps $t_i : C' \rightarrow C_i$, $t_j : C' \rightarrow C_j$. The relations on P are defined by taking $t(p_{i_1}, p_{j_1}) = t'(p_{i_2}, p_{j_2})$ when the corresponding terms in A_{sep}^* and B_{sep}^* are equal, i.e., $t_{i_1}(a_{i_1}) = t'_{i_2}(a_{i_2})$ and $t_{j_1}(b_{j_1}) = t'_{j_2}(b_{j_2})$.

This presheaf P is clearly finitely generated, and we have a diagram

$$\begin{array}{ccc} P & \xrightarrow{\pi_2} & B_{\text{sep}}^* \\ \pi_1 \downarrow & & \downarrow \beta^* \\ A_{\text{sep}}^* & \xrightarrow{\alpha^*} & \text{Hom}(-, C)_{\text{sep}} \end{array}$$

This diagram is not in general a pullback diagram, but it suffices to show that it is mapped to one by the localization functor.

Let X be any functor, and suppose we have a commutative diagram

$$\begin{array}{ccc} aX & \xrightarrow{x_2} & B \\ x_1 \downarrow & & \downarrow \beta \\ A & \xrightarrow{\alpha} & a\text{Hom}(-, C) \end{array}$$

The map $x_1 : aX \rightarrow A$ is given by a map $x'_1 : S \rightarrow A'_{\text{sep}}$, where S is a dense subobject of X . The sep-functor preserves monomorphisms, so A_{sep}^* is a subobject of A'_{sep} , and since both objects are mapped to aB by the associated sheaf functor, it is a dense subobject. The pullback of A_{sep}^* along x'_1 is also a dense subobject of X , so there is a map $x_1^* : S \rightarrow A_{\text{sep}}^*$ with $a(x_1^*) = x_1$. Similarly, we may assume there is a map $x_2^* : S \rightarrow B_{\text{sep}}^*$ with $a(x_2^*) = x_2$ (we can assume both maps have the same domain by taking the intersection of the two domains).

Suppose S is generated by elements s_k each of sort S_k . The maps x_1^*, x_2^* send each of these generators s_k to $x_1^*(s_k) \in A_{\text{sep}}^* S_k$, $x_2^*(s_k) \in B_{\text{sep}}^* S_k$. This pair is represented by an element $(x_1^*(s_k), x_2^*(s_k)) \in PS_k$. Defining this on each generator gives us a transformation $\tilde{x} : S \rightarrow P$. We observe that this is indeed a transformation since any relations that are required to hold in P hold in each of its two components, by the assumption that x_1^* and x_2^* were transformations themselves.

We still need to show that the factorization is unique. Suppose there is another transformation $x' : S' \rightarrow P$, such that S' is dense in X , and $a(\pi_1 x') = x_1$ and $a(\pi_2 x') = x_2$. Then by Lemma 3.8, there is a subobject of T of S and S' on which $\pi_1 x'$ agrees with \tilde{x}_1 , and $\pi_2 x'$ agrees with \tilde{x}_2 ; it follows by the definition of \tilde{x} that $\tilde{x}|_T = x'|_T$. But if this is the case then x' and \tilde{x} represent the same transformation $aX \rightarrow aP$. This concludes the proof. \square

Proof of theorem 3.7: The topos of sheaves on a site (\mathcal{C}, J) always has the representable functors $a\text{Hom}(-, C)$ as a generating set; we have just shown that these will be coherent. The converse follows immediately from the Comparison Lemma, 3.1 and Lemma 3.6.

3.3 Modules over a sheaf of rings

Let \mathcal{E} be a topos of presheaves, i.e., $\mathcal{E} = \mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ for some small category \mathcal{C} . A ring object in \mathcal{E} is a *presheaf of rings* on \mathcal{C} - this is a presheaf $R : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Sets}$ such that RC has a ring structure for every object C in \mathcal{C} , and for each map $f : C \rightarrow C'$ in \mathcal{C} , the map $Rf : RC' \rightarrow RC$ is a morphism of rings. We write $\mathbf{Rings}(\mathcal{E})$ for the category of ring objects in \mathcal{E} .

If J is a topology on \mathcal{C} , then a ring object in $\mathbf{Sh}(\mathcal{C}, J)$ is a presheaf of rings such that the underlying presheaf of sets is a sheaf.

In particular, since the localization functor $a : \mathbf{Sets}^{\mathcal{C}^{\text{op}}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$ preserves finite products, we see that the localization of a presheaf of rings is also a sheaf of rings, and this is a reflection functor from $\mathbf{Rings}(\mathbf{Sets}^{\mathcal{C}^{\text{op}}}) \rightarrow \mathbf{Rings}(\mathbf{Sh}(\mathcal{C}, J))$ (and more generally, if $f : \mathcal{E} \rightarrow \mathcal{F}$ is a geometric morphism of toposes, this defines a ‘geometric morphism’ $\mathbf{Rings}(\mathcal{F}) \rightarrow \mathbf{Rings}(\mathcal{E})$, that is, an adjoint pair of functor between these two categories with the left adjoint preserving finite limits).

If $(R, 0_R, 1_R, -_R, +_R, \times_R)$ is a ring object in a topos, we define a *right R -module object* $(M, 0_M, -_M, +_M, \times_M)$ in \mathcal{C} to be an abelian group object $(M, 0_M, -_M, +_M)$ together with a map $\times_M : M \times R \rightarrow M$ satisfying the commutativity conditions required by modules; for example, to show the multiplication is distributive over

addition, we stipulate that the following diagram must commute:

$$\begin{array}{ccc}
 M \times R \times R & \xrightarrow{(\times)_M \times \text{id}_R} & M \times R \\
 \text{id}_M \times (\times_R) \downarrow & & \downarrow \times_M \\
 M \times R & \xrightarrow{\times_M} & M
 \end{array}$$

Morphisms of R -module objects are defined similarly.

If R is a presheaf of rings over some small category \mathcal{C} , then an R -module object M in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ is a ‘presheaf of R -modules’ - for each object C in \mathcal{C} , MC will be an RC -module, and for a map $f : C \rightarrow C'$, the map $Mf : MC' \rightarrow MC$ will be an RC' -linear map, where MC is considered with the action of RC' on it defined by the map $Rf : RC' \rightarrow RC$. We denote the category of presheaves of R -modules over a presheaf of rings by $\mathbf{PreMod}\text{-}R$. If R is a sheaf of rings we denote the category of sheaves of R -modules by $\mathbf{Mod}\text{-}R$.

In particular, we see that since the localization functor $a : \mathbf{Sets}^{\mathcal{C}^{\text{op}}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$ preserves finite limits, if M is a presheaf of R -modules for some presheaf of rings R on a small category \mathcal{C} , then aM will be a sheaf of aR -modules.

Now let N be a sheaf of aR -modules. The presheaf of rings R has an action on N , given by

$$N \times R \xrightarrow{\eta_R \times \text{id}_N} N \times aR \xrightarrow{\times_N} N$$

and one can easily show that N is a presheaf of R -modules with this action.

Theorem 3.10. *Let (\mathcal{C}, J) be a site, with associated sheaf functor $a : \mathbf{Sh}(\mathcal{C}, J) \rightarrow \mathbf{Sets}^{\mathcal{C}^{\text{op}}}$, and inclusion functor $i : \mathbf{Sets}^{\mathcal{C}^{\text{op}}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$. Suppose R is a presheaf of rings on \mathcal{C} . Then a and i induce functors $a' : \mathbf{PreMod}\text{-}R \rightarrow \mathbf{Mod}\text{-}aR$, $i' : \mathbf{Mod}\text{-}aR \rightarrow \mathbf{PreMod}\text{-}R$, and this expresses $\mathbf{Mod}\text{-}aR$ as a localization of the category $\mathbf{PreMod}\text{-}R$.*

Proof. We have already described the functors a' and i' , and a' preserves finite limits since it commutes with the forgetful functor. It remains to show that a' is left adjoint to i' , or equivalently, that given a presheaf of R -modules M and a sheaf of aR modules N , then a map $f : M \rightarrow N$ in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ is a morphism in $\mathbf{PreMod}\text{-}R$ if and only if

the corresponding map $af : aM \rightarrow N$ in $\mathbf{Sh}(\mathcal{C}, J)$ is a morphism in $\mathbf{Mod}\text{-}aR$. We show that f commutes with the multiplication by R if and only if af commutes with the multiplication by aR ; the proof that f is an abelian group map if and only if af is is similar.

Consider the diagram below, where $f = af \cdot \eta_M$:

$$\begin{array}{ccccc}
 M \times R & \xrightarrow{f \times \text{id}_R} & N \times R & & \\
 \downarrow \times_M & \searrow \eta_{(M \times R)} & \downarrow \text{id}_N \times \eta_R = \eta_{(N \times R)} & & \\
 & aM \times aR & \xrightarrow{af \times \text{id}_{aR}} & N \times aR & \\
 & \downarrow a(\times_M) & & \downarrow \times_N & \\
 M & \xrightarrow{\eta_M} & aM & \xrightarrow{af} & N
 \end{array}$$

Our claim is that the inner square commutes if and only if the outer square does. To see this, notice that the inner square is the image of the outer square under the functor a ; thus if the outer square commutes, the inner square must commute also. Now suppose the inner square commutes. Then $\times_N \cdot \eta_{(N \times R)} \cdot (f \times \text{id}_R) = \times_N \cdot (af \times \text{id}_{aR}) \cdot \eta_{(M \times R)} = af \cdot a(\times_M) \cdot \eta_{M \times R} = af \cdot \eta_M \cdot \times_M$, and thus the outer square commutes, as required. \square

In the paper of Prest and Ralph, [24], the following result was shown: let X be a topological space with a basis of compact open sets, and let R be a sheaf of rings on X (we refer to a topological space with a sheaf of rings as a *ringed space*). Then the category of R -modules is locally finitely presentable.

In that paper, it was asked under what conditions this result generalizes to an arbitrary Grothendieck topos. This question can be answered using the characterizations given up to this point.

Theorem 3.11. *If \mathcal{E} is a locally finitely presentable topos (respectively locally finitely generated) and R is a ring object in \mathcal{E} , then $\text{Mod}\text{-}R$, the category of R -module objects in \mathcal{E} , is locally finitely presentable (respectively locally finitely generated).*

Proof. Let \mathcal{E} be a locally finitely presentable topos. Then by Proposition 3.4, \mathcal{E} is equivalent to $\mathbf{Sh}(\mathcal{C}, J)$, where J is a topology on a small category \mathcal{C} with the property

that every covering sieve on an object C in \mathcal{C} contains a finitely presentable covering sieve.

We use the essentially algebraic theory describing the objects in $\mathbf{Sh}(\mathcal{C}, J)$ described in section 3.1. Sheaves of R -modules can be described by adding more functions and equations to this theory.

To define the category of R -module objects in $\mathbf{Sh}(\mathcal{C}, J)$, we add total operation symbols to the above signature. For a given sort \mathbf{C} (that is, each object C in \mathcal{C}), we add a function symbol $\mathbf{r} : \mathbf{C} \rightarrow \mathbf{C}$, for each element $r \in RC$. We also add a constant symbol $0_{\mathbf{C}}$ of sort \mathbf{C} , and function symbols $+_{\mathbf{C}} : \mathbf{C} \times \mathbf{C} \rightarrow \mathbf{C}$ and $-_{\mathbf{C}} : \mathbf{C} \rightarrow \mathbf{C}$. We expand E to include equations stating that with the operations so defined, the collection of elements of the sort \mathbf{C} is an RC -module for each object C , and each function symbol $\mathbf{f} : \mathbf{B} \rightarrow \mathbf{C}$ is an RB -linear map.

This gives us a description of $\text{Mod-}R$ as a finitary essentially algebraic category; it is therefore locally finitely presentable.

To get the corresponding result for locally finitely generated toposes, we use a similar argument, but the sets $\text{Def}(\sigma)$ are allowed to contain infinitely many equations. The rest of the argument is unchanged. \square

It is well-known that for a topological space X , the category of sheaves on X is locally finitely presentable if and only if the space has a basis \mathcal{B} of compact open sets, see e.g., [17, D3.3.14]. This can be seen as a consequence of [20, II.2.3], which states that a sheaf on the lattice $\text{Op}(X)$ is equivalent to a sheaf on the sublattice consisting of the basis elements \mathcal{B} . If this basis consists of compact open sets, then the essentially algebraic theory described in section 3.1 will be finitary.

Remark. It has been pointed out to us that results in section D5 of [17] can be used to provide a straightforward proof that local λ -presentability of a topos \mathcal{E} implies local λ -presentability of any category of modules in \mathcal{E} . Our proof has a much more model theoretic flavour, and we hope it will be more straightforward to those with experience in this field.

We now turn our attention to the question of when the category of modules on a

ringed space is locally coherent. It should be noted at this point that the definition of coherence we are using is distinct from (and a lot stronger than) the definition of a coherent sheaf of modules used by algebraic geometers, in for example [15, II.5].

Let X be a topological space, and let R be a presheaf of rings on X . Denote by $\text{Op}(X)$ the lattice of open sets of X . Denote by Γ_R the essentially algebraic theory of presheaves of modules over R .

For each open set $U \in \text{Op}(X)$, define a presheaf $R_U : \text{Op}(X)^{\text{op}} \rightarrow \mathbf{Sets}$ by

$$R_U(V) = R(V) \text{ if } V \subseteq U; 0 \text{ if } V \not\subseteq U.$$

Each presheaf-of-modules R_U is the free model of the theory Γ_R generated by a single element of sort U . It follows that the collection $\{R_U \mid U \in \text{Op}(X)\}$ is a generating set of objects for $\mathbf{PreMod}\text{-}R$ - if $\alpha \neq \beta : F \rightarrow G$ are distinct maps in $\mathbf{PreMod}\text{-}R$, then there is some $U \in \text{Op}(X)$ and some $x \in FU$ such that $\alpha_U(x) \neq \beta_U(x)$. The element x represents a map $\tilde{x} : R_U \rightarrow F$, and the inequality $\alpha\tilde{x} \neq \beta\tilde{x}$ holds.

Every presheaf R_U is finitely generated, and is a subobject of the presheaf R , considered as an object in $\mathbf{PreMod}\text{-}R$. In particular, we see as a consequence of Lemma 2.5:

Lemma 3.12. *Let R be a presheaf of rings on a space X . The category $\mathbf{PreMod}\text{-}R$ is locally coherent if and only if R , considered as a sheaf of modules, is coherent as an object in the category.*

Proof. We have just seen that if R is coherent as a module over itself, the objects R_U form a generating set of coherent objects. Now suppose the category $\mathbf{PreMod}\text{-}R$ is locally coherent; then every finitely presentable object is coherent, and since R is finitely presentable in $\mathbf{PreMod}\text{-}R$, it must be coherent. \square

We look for conditions under which a presheaf of rings R is coherent.

Lemma 3.13. *If R is a coherent presheaf of rings on a space X , then for every open set U in X , RU is a coherent ring.*

Proof. Let R be a presheaf of rings on a space, and let U be an open set in X such that RU is not a coherent ring. Then there is some finitely generated ideal $\langle x_1, \dots, x_n \rangle = I \subseteq RU$, such that I is not finitely presentable. But I defines a finitely generated subobject of R - just take the subobject of R generated by x_1, \dots, x_n . This is not finitely presentable; if I had a presentation with only finitely many relations, then taking the relations which occur between terms of sort U would give a finite presentation of IU . \square

A version of the next result appeared as Theorem 2.18 in [25]. In this paper, the result was proved for an arbitrary presheaf of rings R , as long as RU was coherent for each open set $U \in \text{Op}(X)$. As it turns out, the argument in [25] contains a mistake: if F is a finitely presentable module and G a finitely generated submodule, then the sheaf-of-rings structure of R can force G to have infinitely many relations. The argument can be made to work if we insist that the presheaf-of-rings R be finitely presentable, however.

Theorem 3.14. (*cf [25, 2.18]*) *Let R be a presheaf of rings on the space X . If R is finitely presentable, and for each open set $U \in \text{Op}(X)$, RU is a coherent ring, then the presheaf of rings R is coherent as an object in $\mathbf{PreMod}\text{-}R$.*

Proof. Let $\langle x_1, \dots, x_n \mid r_1(\mathbf{x}), \dots, r_m(\mathbf{x}) \rangle$ be a presentation of R in the language of presheaves of rings over \mathcal{C} . Each element x_i has sort U_i where U_i is some open set in $\text{Op}(X)$. Suppose for some open set U_i , we have an open set $U \subseteq U_i$; then denote by x_i^U the restriction of the element x_i to U (i.e., the image of x_i under the restriction map $RU_i \rightarrow RU$). Each relation $r_j(\mathbf{x})$ is an equation between terms in the variables $\mathbf{x} = x_1, \dots, x_n$, and each equation is between terms in some sort V_j , where $V_j \in \text{Op}(X)$.

On any open set $U \in \text{Op}(X)$, the ring RU is generated by the set of elements x_i^U , for those elements x_i where $U \subseteq U_i$. The relations that hold on these elements $\{x_i^U \mid U \subseteq U_i\}$ are precisely those induced by those relations $r_j(\mathbf{x})$ which have $U \subseteq U_j$. In particular, if U and U' are two open sets that are contained in precisely the same open sets U_i and V_j (that is, for every $i = 1, \dots, n$, $U \subseteq U_i$ if and only if $U' \subseteq U_i$,

and for every $j = 1, \dots, m$, $U \subseteq U_j$ if and only if $U' \subseteq U_j$) then the rings RU and RU' are isomorphic (since they have the same presentation).

To prove that the category $\mathbf{PreMod}\text{-}R$ is locally coherent, it suffices to prove that the presheaf-of-modules R is a coherent object of this category. Since $\mathbf{PreMod}\text{-}R$ is an abelian category, it suffices to show that any finitely generated subobject of R is finitely presentable. To see this, suppose that $I \subseteq R$ is a finitely generated subobject of R ; let y_1, \dots, y_l be a generating set of elements, where each element y_k is of sort W_k . Suppose we have two open sets $U \subseteq U'$ that are contained in precisely the same open sets from the collection $U_1, \dots, U_n, V_1, \dots, V_m, W_1, \dots, W_l$. Then RU and RU' are isomorphic, and this isomorphism restricts to an isomorphism $IU \cong IU'$, since IU will be the subobject of RU generated by the elements y_k^U for each k with $U \subseteq W_k$, and IU' will be the generated by the elements $y_k^{U'}$ for precisely the same values of k from $1, \dots, l$.

Thus the presheaf-of-rings I is completely described by its presentation on the open sets which are intersections of subsets of the set $\{U_1, \dots, U_n, V_1, \dots, V_m, W_1, \dots, W_l\}$. There are only finitely many such intersections, and since RU is coherent on every open set U , IU is finitely presentable on each open set U that is such an intersection. Combining the finite presentation of I on each intersection from this set, we can write down a finite presentation for the whole presheaf I . \square

The next result gives us a condition for categories of sheaves of modules (as opposed to just presheaves) to be locally coherent.

Theorem 3.15. *Let (\mathcal{C}, J) be a site such that the inclusion functor $i : \mathbf{Sh}(\mathcal{C}, J) \rightarrow \mathbf{Sets}^{\mathcal{C}^{op}}$ preserves directed colimits, and let R be a presheaf of modules on \mathcal{C} such that $\mathbf{PreMod}\text{-}R$ is locally coherent. Then the category $\mathbf{Mod}\text{-}aR$ is locally coherent also.*

Proof. Suppose (I, \leq) is some directed partial order, and we are given a directed diagram in $\mathbf{Mod}\text{-}aR$,

$$\{D_i \xrightarrow{d_{ij}} D_j \mid i \leq j \in (I, \leq)\}.$$

This is a directed diagram in $\mathbf{PreMod}\text{-}R$ also, and the directed colimit of this diagram exists in $\mathbf{PreMod}\text{-}R$, and has as its underlying presheaf of sets the colimit


of the underlying presheaves of sets (see [1, 3.4(4), 3.6(6)]). Let L be the directed colimit in $\mathbf{PreMod}\text{-}R$, so we have an action of R on L denoted $\times_L : L \times R \rightarrow L$.

Since the inclusion functor $i : \mathbf{Sh}(\mathcal{C}, J) \rightarrow \mathbf{Sets}^{c^{op}}$ preserves directed colimits, L is a sheaf, and the map $a(\times_L) : L \times aR \rightarrow L$ is an action of aR on L , with respect to which L is the colimit of the diagram in $\mathbf{Mod}\text{-}aR$. \square

It is nevertheless possible to find a ring object R in a locally coherent topos, such that the category of R -modules is not locally coherent, as demonstrated by the next result.

Theorem 3.16. *Let X be a topological space with infinitely many open sets. Then there is a sheaf of rings R on X , with the property that RU is a coherent ring for every open set U in X , but the category $\mathbf{Mod}\text{-}R$ is not locally coherent.*

Proof. If X has infinitely many open sets, then the distributive lattice $\text{Op}(X)$ has infinitely many elements. Recall that such a lattice must contain an infinite chain. For suppose it does not. Then the lattice has finite height n . Since the lattice has infinitely many elements, there is a least height $i < n$ such that $\text{Op}(X)$ has infinitely many elements of height i . Note that since the empty set is the unique element of height 0, $i > 0$. There are only finitely many elements of height $i - 1$, so there must be one element, U say, with infinitely many elements of height i above U . Choose a countable collection of these, $\{U_i\}_{i \in \mathbb{N}_0}$.



For a given i , consider whether U_i is contained in the union of the other elements $\bigvee_{j \neq i} U_j$. Only finitely many of the U_i can have this property - if infinitely many did, then we could construct an infinite chain from taking their finite unions. So there exists at least one open set which does not - U_0 , say. So we may assume $U_0 \subseteq \bigvee_{i \geq 1} U_i$.

For each $i \geq 1$, $U_0 \cap U_i = U$. So we have the equalities

$$\bigvee_{i \geq 1} (U_0 \cap U_i) = U$$

and

$$U_0 \cap \bigvee_{i \geq 1} U_i = U_0.$$

Together, these equalities contradict the fact that $\text{Op}(X)$ is a Heyting algebra. So $\text{Op}(X)$ must have an infinite chain.

Now suppose the lattice $\text{Op}(X)$ contains an infinite chain. Then in particular, it contains a chain isomorphic to ω , or a chain isomorphic to ω^{op} . We deal with these two cases separately.

First of all, we look at the case when $\text{Op}(X)$ contains a sequence of open sets isomorphic to ω , say

$$U_0 \longrightarrow U_1 \longrightarrow \dots \quad U_n \longrightarrow \dots \longrightarrow U_\infty$$

We may assume that U_0 is the empty set, that $U_\infty = \bigcup_{i=0}^{\infty} U_i$, and that every open set U_α is connected (if U_α is the smallest disconnected set, replace it with the connected component of U_α that contains $U_{\alpha-1}$).

In this case, for each U_i , we define $RU_i = \mathbb{Z}/2^i\mathbb{Z}$. We define RU_∞ to be \mathbb{Z}_2 , the ring of 2-adic integers. For an arbitrary connected open set $V \subseteq X$, there is a smallest number $i_V \in \mathbb{N} \cup \{\infty\}$ such that $V \subseteq U_{i_V}$. Define $RV = RU_{i_V}$. For a disconnected subset V , define RV to be the product over the connected components.

The presheaf R so described is a sheaf. It suffices to check the sheaf condition on connected subsets V of X , since R sends disjoint unions to the appropriate product by definition.

Suppose V is a connected subset of X , and let $\{V_j\}_{j \in J}$ be a cover of V , and take a matching family $x_j \in RV_j$ of elements of R . Suppose $RV = RU_\alpha$ (that is, U_α is the smallest set in the chain such that $V \subseteq U_\alpha$). But then there is some point $p \in V$ with $p \notin U_i$ for any $i < \alpha$. But the point p is in V_j for some $j \in J$, and in

particular $RV_j = RU_\alpha = RV$. Now the amalgamation for the matching family will be the element $x_j \in RV_j = RV$.

The sheaf R is finitely presentable as a module over itself. As before, we consider the subsheaf M generated by $2 \in RX$. On each open set V in X , MV consists of those elements of RV which admit a division by two. The same argument as for R shows that this is a sheaf.

We need to show that M is not finitely presentable. The argument above shows that if V is covered by subsets V_j then $RV_j = RV$ for some j , and consequently $MV_j = MV$ also. Thus if the presentation includes a relation on some open subset V , this cannot be used to derive relations on subsets strictly containing V . Now suppose there is a finite presentation for M . There is some $i \in \mathbb{N}$ such that the finite presentation does not induce relations on MU_i . But MU_i is not free. So there cannot be any finite presentation for M .

Now we examine the case when $\text{Op}(X)$ contains an infinite chain isomorphic to ω^{op} , say

$$0 = U_\infty \longrightarrow \dots \longrightarrow U_n \longrightarrow \dots \longrightarrow U_1 \longrightarrow U_0 = X.$$

Let k be a field. Define a sheaf of rings R on X as follows: let $RX = k[x_0]$. For an open set V that is not contained in any U_i for $i \in \mathbb{N}$, set $RV = \mathbf{0}$, the one element ring. For any other open set $V \subseteq U$, let n be the smallest natural number such that $V \subseteq U_n$ (if V is contained in every U_n , set $n = \infty$). Then define

$$RV = k[x_0, x_1, \dots, x_n] / \langle x_0 x_i = x_0 \rangle_{1 \leq i \leq n}.$$

The restriction maps are the canonical inclusions of the polynomial rings. Every ring RV is coherent (since we take a quotient of a finite polynomial ring by a finitely generated ideal).

We can show that R is a sheaf in a similar way to the first case: if V is covered by a collection of open sets V_i , there is some V_i with $RV_i = RV$, and this is how we find an amalgamation whenever we have a matching family of elements of R for the cover.

Now let M be the subobject of R in $\text{Mod-}R$ generated by the object x_0 of sort X . For each open set V of X with $RV = k[x_0, \dots, x_n]/\langle x_0x_i = x_0 \rangle_{1 \leq i \leq n}$, MV is given by the presentation

$$MV = \langle y | yx_i = y \rangle_{1 \leq i \leq n}.$$

The same argument as for R shows that M is a sheaf. The object M is finitely generated in $\text{Mod-}R$ (by x_0), but a finite presentation for M would only mention finitely many of the variables x_i in M , so there would be some x_k that is not mentioned in the presentation. But then MU_k would not satisfy $yx_k = y$. So M is not finitely presentable. \square

Chapter 4

Krull-Gabriel Dimension

4.1 Supercompactness and regularity

In this section we introduce the idea of supercompact and regular objects in a topos. The idea of a supercompact object is analogous to that of a finitely generated object, only instead of capturing the idea of an object having a finite set of generators, the idea of supercompactness is to describe those objects with only a single generator. This gives rise to a notion of local supercompactness, and morphisms of toposes which preserve this structure. In particular, localizations of a classifying topos which preserve this structure correspond to theories in regular logic, just as finite-type localizations correspond to theories in a coherent logic. Supercompact and regular objects were introduced in [17, D3.1.12]. We describe these now.

An object C in a topos \mathcal{E} is said to be *supercompact* if whenever there is an epimorphic family of morphisms $\{f : D_f \rightarrow C \mid f \in I\}$, there is (at least) one morphism $f : D_f \rightarrow C$ in the family that is itself an epimorphism. We say C is *regular* if it is supercompact and for any pullback diagram

$$\begin{array}{ccc} B \times_C B' & \longrightarrow & B' \\ \downarrow & & \downarrow \\ B & \longrightarrow & C \end{array}$$

where B and B' are supercompact, the object $B \times_C B'$ is supercompact also. The

topos \mathcal{E} is *locally supercompact* if it is cocomplete and has a strong generating set of supercompact objects, and it is *locally regular* if it is cocomplete and has a strong generating set of regular objects.

The definition of supercompactness can be rephrased in a manner similar to the definition of finite presentation.

Lemma 4.1. *An object C in a topos \mathcal{E} is supercompact if and only if whenever there is a epimorphic family $\{m_i : B_i \rightarrow B\}_{i \in I}$ of morphisms in \mathcal{E} , such that each m_i is a monomorphism, then every map $f : C \rightarrow B$ factors through one of the maps m_i , that is, there is a map $f' : C \rightarrow B_i$ with $f = m_i f'$.*

Proof. Let C be a supercompact object in \mathcal{E} . Suppose we have an epimorphic family $\{m_i : B_i \rightarrow B\}_{i \in I}$ and a map $f : C \rightarrow B$ as described. Then the collection of maps $f^{-1}(m_i) : f^{-1}(B_i) \rightarrow C$ defined by pulling back the maps m_i along f is an epimorphic family of monics. In particular, since C is supercompact, there is some $j \in I$ such that $f^{-1}(m_j)$ is an epimorphism. Since $f^{-1}(m_j)$ is monic and epic, and the topos \mathcal{E} is a balanced category, $f^{-1}(m_j)$ is an isomorphism, and the map $f^{-1}(B_j) \rightarrow B_j$ in the pullback square for m_j is the factorization required.

Now suppose C has the property stated. Suppose there is an epimorphic family $\{f_i : C_i \rightarrow C\}_{i \in I}$ in \mathcal{E} . For each f_i , let $m_i : \text{Im}(C_i) \rightarrow C$ be the image factorization of f_i (this exists by [20, IV.6.1]). The maps m_i form an epimorphic family, and so the identity map $1_C : C \rightarrow C$ factors through one of these maps, say $1_C = m_j p$. The map m_j is both monic and epic, and so is an isomorphism. It follows that the map f_j is an epimorphism. \square

It follows from this characterization that supercompact objects in a topos are finitely generated, since a directed union is in particular an epimorphic family of monics.

There is an analogue of Lemmas 2.1 and 2.2 for supercompact objects.

Lemma 4.2. *Let \mathcal{E} be a locally supercompact topos, and let \mathcal{G} be a generating set of supercompact objects. Then an object C in \mathcal{E} is supercompact if and only if there is an epimorphism $G \rightarrow C$, with $G \in \mathcal{G}$.*

Proof. The collection of all mappings $\{G_i \rightarrow C \mid G_i \in \mathcal{G}\}$ is an epimorphic family, so if C is supercompact, there must be an epimorphism $G_i \rightarrow C$ for some $G_i \in \mathcal{G}$, by supercompactness.

Conversely, suppose we have an object C with an epimorphism $g : G \rightarrow C$, with G supercompact, and let $\{C_i \xrightarrow{c_i} C\}$ be an epimorphic family. Then pulling back, we have that $\{g^*C_i \xrightarrow{g^*c_i} G\}$ is an epimorphic family, and therefore for some map c_i in the family, $g^*c_i : g^*C_i \rightarrow G$ is an epimorphism. It follows that c_i is an epimorphism, since the epimorphism $g.g^*c_i$ factors through it. \square

Now suppose \mathcal{C} is a locally finitely presentable category. Then \mathcal{C} is the category of models for a Cartesian theory \mathbb{T} . The classifying topos for \mathbb{T} is the functor category $(\mathcal{C}_{\mathbb{T}}, \mathbf{Sets})$, where the objects of $\mathcal{C}_{\mathbb{T}}$ are (equivalence classes of) Cartesian formulas over \mathbb{T} , and the morphisms are (equivalence classes of) functions whose graphs are definable by coherent formulas.

There is a way of associating the finitely presentable \mathbb{T} -models with Cartesian formulas, and vice versa.

If the formula $\phi(\mathbf{x})$ is a finite conjunction of atomic formulas, then there is an object $A \in \mathbf{fp}(\mathcal{C})$ and a tuple of generators $\mathbf{a} \in \phi(A)$ with the property that if B is any object and $\mathbf{b} \in \phi(B)$ then the map $\mathbf{a} \mapsto \mathbf{b}$ extends to a homomorphism $A \rightarrow B$. We say that the pair (A, \mathbf{a}) is a *free realization* of the formula ϕ .

If $C \in \mathbf{fp}(\mathcal{C})$ and \mathbf{c} is a tuple of generators for C , then there is a finite conjunction of atomic formulas $\phi(\mathbf{x})$ with $\mathbf{c} \in \phi(C)$ such that (C, \mathbf{c}) is a free realization of ϕ , we say ϕ is a *presentation formula* for C .

Let $\phi(\mathbf{x})$ be an arbitrary formula that is Cartesian with respect to \mathbb{T} . Suppose $\phi(\mathbf{x})$ has the form $\exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y})$, where ψ is a conjunction of atomic formulas. Then the formula $\rho(\mathbf{x}, \mathbf{x}', \mathbf{y})$ given by:

$$\phi(\mathbf{x}) \wedge \psi(\mathbf{x}', \mathbf{y}) \wedge (\mathbf{x} = \mathbf{x}')$$

defines an isomorphism $[\phi] \rightarrow [\psi]$ in $\mathcal{C}_{\mathbb{T}}$. It follows that for every Cartesian formula ϕ over \mathbb{T} , $[\phi]$ is isomorphic in $\mathcal{C}_{\mathbb{T}}$ to the equivalence class of a presentation formula.

Sending each quantifier-free formula $\phi(\mathbf{x})$ to a free realization, we can show that the classifying topos $(\mathcal{C}_{\mathbb{T}}, \mathbf{Sets})$ is equivalent to the functor category $(\mathbf{fp} \mathcal{C}, \mathbf{Sets})$. Recall from section 2.6 that a geometric formula $\phi(\mathbf{x})$ over the signature of \mathbb{T} gives rise to a functor $\mathbf{Mod}(\mathbb{T}) \rightarrow \mathbf{Sets}$, also denoted ϕ . If C is a finitely presented object in \mathcal{C} with presentation formula $\phi(\mathbf{x})$, we have just seen that for any other object B in \mathcal{C} , maps $C \rightarrow B$ correspond to tuples in $\phi(B)$. This sets up an isomorphism between the representable functor $\mathrm{Hom}(C, -)$ and the functor $\phi : \mathbf{Mod}(\mathbb{T}) \rightarrow \mathbf{Sets}$.

Let F be a functor $F : \mathbf{fp}(\mathcal{C}) \rightarrow \mathbf{Sets}$. A *functorial equivalence relation* on F is a parallel pair of morphisms $(r_1, r_2) : E \rightrightarrows F$ that defines an equivalence relation on F in the functor category $(\mathcal{C}, \mathbf{Sets})$.

It follows from the Yoneda lemma that all the representable presheaves in $(\mathbf{fp}(\mathcal{C}), \mathbf{Sets})$ are supercompact - if $\mathrm{Hom}(-, C)$ is a representable functor and there is a strong epimorphic family $\{F_i \rightarrow \mathrm{Hom}(-, C) \mid i \in I\}$ in $(\mathbf{fp} \mathcal{C}, \mathbf{Sets})$, then in particular the identity map $1_C : C \rightarrow C$ must lie in the image of F_j for some $j \in I$. But then the map $F_j \rightarrow \mathrm{Hom}(-, C)$ is epimorphic. The representable presheaves form a strong generating set.

In particular, we get the following result, which follows from observing the proofs of 3.2.2, 3.2.3 and 3.2.5 in [26].

Proposition 4.3. *An object F in $(\mathbf{fp} \mathcal{C}, \mathbf{Sets})$ is supercompact if and only if it is isomorphic to the functor ϕ/E , where ϕ is a quantifier free regular formula, and E is a functorial equivalence relation.*

Proof. Let (A, \mathbf{a}) be a free realization of $\phi(\mathbf{x})$; then the functor ϕ is representable, and therefore supercompact. Since ϕ/E is a quotient of ϕ , it is supercompact also, by Lemma 4.2.

The representable functors form a strong generating set for $(\mathbf{fp} \mathcal{C}, \mathbf{Sets})$, so by Lemma 4.2, there must be an epimorphism $e : \mathrm{Hom}(C, -) \rightarrow F$ for some finitely presented object C in \mathcal{C} . Let ϕ be a presentation formula for C , so $\mathrm{Hom}(C, -) \cong \phi$.

Let $R \subseteq \phi \times \phi$ be the functor given by

$$M \mapsto \{(\mathbf{x}, \mathbf{y}) \in \phi(M)^2 : e_M(\mathbf{x}) = e_M(\mathbf{y})\}$$

and let $E(\mathbf{x}, \mathbf{y})$ be the equivalence relation generated by R . Then $F \cong \phi/E$ as required. \square

Similarly, we extract from the proof of [26, 3.2.7]:

Lemma 4.4. *For every morphism $f : \phi/E_1 \rightarrow \psi/E_2$ between supercompact objects in this category, there is a regular formula ρ such that for each model M , $\rho(M)$ is the graph of the map $f_M : \phi/E_1(M) \rightarrow \psi/E_2(M)$.*

Proof. Let $\alpha : \phi/E_1 \rightarrow \psi/E_2$ be such a natural transformation. Let (A, \mathbf{a}) be a free realisation of the formula $\phi(\mathbf{x})$. Choose a representative \mathbf{a}' for the E_2 -equivalence class $\alpha_A[\mathbf{a}]_{E_1}$. Since the tuple \mathbf{a} generates A , there will be a tuple \mathbf{t} of terms such that $\mathbf{t}(\mathbf{a}) = \mathbf{a}'$. Let $\rho(\mathbf{x}, \mathbf{y})$ be the formula

$$\phi(\mathbf{x}) \wedge \mathbf{t}(\mathbf{x}) = \mathbf{y}.$$

Suppose that \mathbf{c}, \mathbf{c}' are tuples from an object C and that $\rho(\mathbf{c}, \mathbf{c}')$ holds. Then there is a map $f : A \rightarrow C$ defined by $\mathbf{a} \mapsto \mathbf{c}$. Since α is a natural transformation, the following diagram commutes, where f_* denotes the obvious induced map.

$$\begin{array}{ccc} \frac{\phi}{E_1}(A) & \xrightarrow{\alpha_A} & \frac{\psi}{E_2}(A) \\ f_* \downarrow & & \downarrow f_* \\ \frac{\phi}{E_1}(C) & \xrightarrow{\alpha_C} & \frac{\psi}{E_2}(C) \end{array}$$

Clearly $\mathbf{c} \in \phi(C)$ so there is an equivalence class $[\mathbf{c}]_{E_1}$ for which the following equations hold:

$$\begin{aligned} \alpha_C[\mathbf{c}]_{E_1} &= \alpha_C[f(\mathbf{a})]_{E_1} = \alpha_C f_*[\mathbf{a}]_{E_1} = f_* \alpha_A[\mathbf{a}]_{E_1} = f_*[\mathbf{a}']_{E_2} \\ &= f_*[\mathbf{t}(\mathbf{a})]_{E_2} = [f(\mathbf{t}(\mathbf{a}))]_{E_2} = [\mathbf{t}(f(\mathbf{a}))]_{E_2} = [\mathbf{t}(\mathbf{c})]_{E_2} = [\mathbf{c}']_{E_2}. \end{aligned}$$

So we see that if $\rho(\mathbf{c}, \mathbf{c}')$ holds in C , then $\alpha_C[\mathbf{c}]_{E_1} = [\mathbf{c}']_{E_2}$. We now check that ρ defines the natural transformation α .

Let $\mathbf{a} \in \phi(A)$. Clearly $\rho(\mathbf{a}, \mathbf{t}(\mathbf{a}))$ holds. This implies that $\alpha_A[\mathbf{a}]_{E_1} = [\mathbf{t}(\mathbf{a})]_{E_2}$, so we must have $\mathbf{t}(\mathbf{a}) \in \psi(A)$. Now suppose that $E_1(\mathbf{a}, \mathbf{a}')$, $\rho(\mathbf{a}, \mathbf{c})$, and $\rho(\mathbf{a}', \mathbf{c}')$ hold. Then $[\mathbf{c}]_{E_2} = \alpha_A[\mathbf{a}]_{E_1} = \alpha_A[\mathbf{a}']_{E_1} = [\mathbf{c}']_{E_2}$. So we have $\mathbf{c}, \mathbf{c}' \in \psi(A)$ and $E_2(\mathbf{c}, \mathbf{c}')$. Hence ρ determines the natural transformation $\alpha : \phi/E_1 \rightarrow \psi/E_2$. \square

Proposition 4.5. *Let F be a presheaf in $(\mathbf{fp} \mathcal{C}, \mathbf{Sets})$. F is regular if and only if it is of the form ϕ/θ , where ϕ and θ are regular formulas and θ defines an equivalence relation on ϕ .*

Proof. Suppose F is regular; then $F \cong \phi/E$, where ϕ is a regular formula and E is a functorial equivalence relation on ϕ . We can construct the following pullback diagram in $(\mathbf{fp}(\mathcal{C}), \mathbf{Sets})$.

$$\begin{array}{ccc} \phi \times_F \phi & \longrightarrow & \phi \\ \downarrow & & \downarrow \\ \phi & \longrightarrow & \phi/E \end{array}$$

Pullbacks in functor categories are defined pointwise, so for each object C in $\mathbf{fp}(\mathcal{C})$, the diagram

$$\begin{array}{ccc} (\phi \times_F \phi)(C) & \longrightarrow & \phi(C) \\ \downarrow & & \downarrow \\ \phi(C) & \longrightarrow & \phi/E(C) \end{array}$$

is a pullback in \mathbf{Sets} . Therefore,

$$(\phi \times_F \phi)(C) = \{(\mathbf{x}, \mathbf{y}) \in \phi(C) \times \phi(C) \mid (x, y) \in E(C)\}.$$

That is, $(\phi \times_F \phi)(C) \cong E(C)$. The isomorphism is natural in C , so we get an isomorphism $\phi \times_F \phi \cong E$. Thus E is a supercompact object in $(\mathbf{fp}(\mathcal{C}), \mathbf{Sets})$ and it is a subobject of the functor $\phi \times \phi$. The functor $\phi \times \phi$ is given by the formula $\phi(\mathbf{x}) \wedge \phi(\mathbf{y})$ in variables (\mathbf{x}, \mathbf{y}) . By Lemmas 4.3 and 4.4, the functor E has the form ψ/E' for some regular formula ψ and some functorial equivalence relation E' on ψ ,

and the inclusion $E \rightarrow \phi \times \phi$ is of the form

$$\psi/E' \xrightarrow{\rho} \phi \wedge \phi$$

for some regular formula ρ . The functor ψ/E' is then equivalent to the functor θ defined by the formula $(\exists \mathbf{y}, \mathbf{y}')\psi(\mathbf{x}) \wedge \rho(\mathbf{x}, \mathbf{y}, \mathbf{y}')$. This is a regular formula, and F is isomorphic to ϕ/θ .

Now suppose that $F \cong \phi/\theta$, where ϕ and θ are regular formulas. Let G be supercompact and $f : G \rightarrow F$ a map. We can write this as

$$\psi/E \xrightarrow{\rho} \phi/\theta$$

for some regular formulas ψ and ρ , and some functorial equivalence relation E on ψ . Let R be the equivalence relation on $\phi \times \phi$ defined by

$$(x, y)R(x', y') \text{ if and only if } E(x, y) \text{ and } E(x', y').$$

Now let $\gamma(x, y)$ be the (positive primitive) formula


$$\exists w, z(\psi(x) \wedge \psi(y) \wedge \rho(x, w) \wedge \rho(y, z) \wedge \theta(w, z)).$$

Then $G \times_F G \cong \gamma/R$, which is supercompact. □


The next result makes the connection between epimorphic families and directed colimits more explicit.

Proposition 4.6. *Let \mathcal{C} be a balanced locally finitely presented category. Then any epimorphic family $\{A_i \xrightarrow{a_i} A\}_{i \in I}$ in \mathcal{C} is contained in a directed colimit cocone.*

Proof. Suppose we are given an epimorphic family as above. For a finite tuple from I , say $\mathbf{i} = (i_1, \dots, i_n)$, let $X_{\mathbf{i}}$ be the limit of the following diagram:



We now define $A_{\mathbf{i}}$ to be the colimit of the diagram



and the map $a_{\mathbf{i}} : A_{\mathbf{i}} \rightarrow A$ is then the factorization of the maps a_{i_j} through the colimit. Similarly, if \mathbf{j} is a tuple all of whose elements are contained in \mathbf{i} , we can define a map $A_{\mathbf{j}} \rightarrow A_{\mathbf{i}}$. This now gives us a filtered system; let $\{A_{\mathbf{j}} \xrightarrow{c_{\mathbf{j}}} C\}$ be the colimit for this system. We note that the ordering on the tuples here is not important, though the number of times each element of I appears is.

Let $\tilde{a} : C \rightarrow A$ be the factorization of the maps $a_{\mathbf{j}}$ through the cocone. Then \tilde{a} is clearly an epimorphism, since the original epimorphic family $\{a_{\mathbf{i}}\}_{\mathbf{i} \in I}$ factors through it.

It is now sufficient to show that \tilde{a} is monic; if this is true then it will be an isomorphism. To show this, it suffices to show that given any pair of maps $G \xrightarrow[x]{y} C$ with $\tilde{a}x = \tilde{a}y$, where G is finitely presented, it must be the case that $x = y$. So suppose we have two such maps. Then since G is finitely presented, there must be some finite tuple \mathbf{i} from I such that x and y both factor through $c_{\mathbf{i}}$, say $x = c_{\mathbf{i}}x'$, $y = c_{\mathbf{i}}y'$. Now

$$a_{\mathbf{i}}x' = \tilde{a}c_{\mathbf{i}}x' = \tilde{a}x = \tilde{a}y = \tilde{a}c_{\mathbf{i}}y' = a_{\mathbf{i}}y'.$$

This means there is a map $G \xrightarrow{[x',y']} X_{(\mathbf{i},\mathbf{i})}$ such that the factorizations through $X_{\mathbf{i}}$ are x' and y' . Now the maps

$$\begin{array}{l}
 G \xrightarrow{x'} A_{\mathbf{i}} \longrightarrow A_{(\mathbf{i},\mathbf{i})} \\
 G \xrightarrow{y'} A_{\mathbf{i}} \longrightarrow A_{(\mathbf{i},\mathbf{i})}
 \end{array}$$

must be equal, since both are given by

$$G \xrightarrow{[x',y']} X_{(\mathbf{i},\mathbf{i})} \longrightarrow A_{(\mathbf{i},\mathbf{i})}.$$

It now follows by composition with $A_{(\mathbf{i},\mathbf{i})} \xrightarrow{c_{(\mathbf{i},\mathbf{i})}} C$ that $x = y$. \square

Remark. In a topos, monomorphisms are preserved by pushouts [20, IV.10.4]. Thus, if \mathcal{C} is a topos, and all the maps $A_i \xrightarrow{a_i} A$ are monic, then the maps in the directed colimit cocone defined in the proof of Theorem 4.6 are all monic too. In particular, a similar argument to that above shows that in a locally supercompact topos, every epimorphic family of monics is contained in a directed union.

Definition 4.7. *Let \mathcal{C} be a small category. A Grothendieck topology J on \mathcal{C} is of regular type if every covering sieve on an object C in \mathcal{C} contains a supercompact sieve.*

Representable presheaves are supercompact in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$: given an epimorphic family $\{f_i : P_i \rightarrow \text{Hom}(-, C)\}_{i \in I}$ there must be some i such that the map $1_C : C \rightarrow C$ is contained in the image of f_i ; this map f_i is then an epimorphism. It follows by Lemma 4.1, a sieve on an object in \mathcal{C} is supercompact if and only if it is generated by a single arrow.

Proposition 4.8. *Let (\mathcal{C}, J) be a site such that the Grothendieck topology J is of regular type. Let $\{F_i \xrightarrow{m_i} F\}$ be an epimorphic family of monomorphisms in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$, where each F_i is a sheaf. Then if F is separated, it is a sheaf.*

Proof. Let $S \xrightarrow{s} \text{Hom}(-, C)$ be a covering sieve of an object C in \mathcal{C} , and suppose we have a map $S \xrightarrow{f} F$. We need to show that there is a unique extension of f to $\text{Hom}(-, C)$, that is, a unique map $\text{Hom}(-, C) \xrightarrow{\tilde{f}} F$ such that $\tilde{f}a = f$.

By assumption, there is a supercompact covering sieve $S' \xrightarrow{s'} \text{Hom}(-, C)$ contained in S , and since $\{m_i\}_{i \in I}$ is an epimorphic family of monics, the map fs' factors through m_i for some i , say $fs' = m_j f'$. Since F_i is a sheaf, and S' is a covering sieve on C , we can extend f' to the whole of $\text{Hom}(-, C)$, say $f' = gs'$, for some $g : \text{Hom}(-, C) \rightarrow F_j$.

Now since F is separated, there is some monic $F \xrightarrow{m} G$ with G a sheaf. We have that mm_jgs and mf are both extensions of mm_jf' to the whole of S ; since such an extension must be unique, it follows that $mm_jgs = mf$, and since m is monic, we get $m_jgs = f$. Thus f_jg is an extension of f to $\text{Hom}(-, C)$.

To show uniqueness, suppose $\text{Hom}(-, C) \xrightarrow{\tilde{f}} F$ is any other map with $\tilde{f}s = f$. Then $m\tilde{f}s = mf = mc_jgs$, and since G is a sheaf, we have that $mm_jg = m\tilde{f}$, by uniqueness of extensions, and again since m is monic, we have $m_jg = \tilde{f}$. \square

Proposition 4.9. *If J is a regular type topology on a small category \mathcal{C} , then the inclusion $i : \mathbf{Sh}(\mathcal{C}, J) \rightarrow \mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ functor preserves epimorphic families of monics.*

Proof. Suppose we are given an epimorphic family of monomorphisms $\{F_i \xrightarrow{m_i} F\}$ in $\mathbf{Sh}(\mathcal{C}, J)$.

Then extend the family of monics, as in Proposition 4.6 (we carry out this extension in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$, not $\mathbf{Sh}(\mathcal{C}, J)$). For each tuple \mathbf{i} from I , the subobject $F_{\mathbf{i}}$ is the union of all the subobjects F_i mentioned in \mathbf{i} . There is then a monic map $F_{\mathbf{i}} \xrightarrow{m_{\mathbf{i}}} F$, and since F is a sheaf, it follows that $F_{\mathbf{i}}$ is separated, and by the previous proposition is therefore also a sheaf.

Now take the colimit of the diagram (in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$). Since this is a directed colimit of sheaves, and since a is a regular type (and therefore finite type) localization, the colimit of this diagram is a sheaf, and is the same as the colimit calculated in $\mathbf{Sh}(\mathcal{C}, J)$. But by Proposition 4.6, this is precisely the cocone $\{F_{\mathbf{i}} \xrightarrow{m_{\mathbf{i}}} F\}$ which we just constructed.

So given any pair of maps $f, g : F \rightarrow X$ in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ with $fm_i = gm_i$ for each i , we have, using the colimit property on each $F_{\mathbf{i}}$, that $fm_{\mathbf{i}} = gm_{\mathbf{i}}$, and therefore, since f and g are both factorizations of the cocone $\{fm_{\mathbf{i}}\}$, we must have that $f = g$. Thus, the m_i form an epimorphic family in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$. \square

To prove the converse, we need the following:

Lemma 4.10. *Let (\mathcal{C}, J) be a site such that the inclusion functor $i : \mathbf{Sh}(\mathcal{C}, J) \rightarrow \mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ preserves epimorphic families of monics. Then the associated sheaf functor $a : \mathbf{Sets}^{\mathcal{C}^{\text{op}}} \rightarrow \mathbf{Sh}(\mathcal{C}, J)$ preserves supercompact objects.*

Proof. Let P be a supercompact object in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$, and suppose we have an epimorphic family of monomorphisms in $\mathbf{Sh}(\mathcal{C}, J)$, say $\{F_i \xrightarrow{m_i} F\}_{i \in I}$, and a map $f : aP \rightarrow F$. Then since the maps m_i form an epimorphic family of monics in $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$, the map $f\eta_P$ factors through one of them, m_j say, so that we have a map $f' : P \rightarrow F_j$ with $m_j f' = f\eta_P$.

$$\begin{array}{ccc} P & \xrightarrow{f'} & F_j \\ \eta_P \downarrow & \nearrow \bar{f} & \downarrow m_j \\ aP & \xrightarrow{f} & F \end{array}$$

Now we get map $\bar{f} : aP \rightarrow F_j$, with $\bar{f}\eta_P = f'$, and consequently $f\eta_P = m_j f' = m_j \bar{f}\eta_P$ and so $m_j \bar{f} = f$ by the universal property of η_P . Thus \bar{f} is a factorization of f through the epimorphic family, and aP is supercompact. \square

Corollary 4.11. *Let (\mathcal{C}, J) be a site. The topology J is of regular type if and only if the inclusion functor $i : \mathbf{Sh}(\mathcal{C}, J) \rightarrow \mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ preserves epimorphic families of monics.*

Proof. One implication is just Proposition 4.9.

To prove the converse, suppose $s : S \rightarrow \text{Hom}(-, C)$ is a covering sieve of an object C in \mathcal{C} . Since $\mathbf{Sets}^{\mathcal{C}^{\text{op}}}$ is locally supercompact, the collection of all maps $G \rightarrow S$ with G supercompact is an epimorphic family and, taking images, this is true if we take all monics $G \rightarrow S$. Now the localization functor a always preserves epimorphic families of monics, so the collection of all monics $aG \rightarrow aS$ is an epimorphic family. But since S is a covering sieve, we have $aS = a\text{Hom}(-, C)$ and, by the previous proposition, $a\text{Hom}(-, C)$ is supercompact in $\mathbf{Sh}(\mathcal{C}, J)$. Therefore there must be some monic $G \rightarrow S$ such that $aG \rightarrow aS$ is an epimorphism. But since it is both epic and monic, it must be an isomorphism. So G is a supercompact subobject of S with $aG \cong a\text{Hom}(-, C)$, i. e. G is a covering sieve of C also. \square

It is now possible to characterize locally supercompact and locally regular toposes in a manner similar to the characterizations of locally finitely generated and coherent toposes given in section 3.2.

If J is a regular type topology on \mathcal{C} , then for every object C in \mathcal{C} , the functor

$a\text{Hom}(-, C)$ is supercompact, by Lemma 4.10. By Lemma 4.2, an object in $\mathbf{Sh}(\mathcal{C}, J)$ is locally supercompact if and only if it is a quotient of a representable functor. Considered as a model of the theory $\Gamma_{\mathbf{Sh}(\mathcal{C}, J)}$, an object is supercompact if and only if it admits a presentation with a single generator. The next result follows from a similar argument to Theorem 3.5.

Theorem 4.12. *Let J be a regular type topology on a category \mathcal{C} . If F is a supercompact sheaf, there is a supercompact presheaf F' with $aF' = F$.*

We can use an argument similar to the proof of Proposition 3.2 to get the following characterization of locally supercompact toposes.

Lemma 4.13. *A topos \mathcal{E} is locally supercompact if and only if it is equivalent to the category $\mathbf{Sh}(\mathcal{C}, J)$ for some site (\mathcal{C}, J) in which every covering sieve in J contains a covering sieve generated by a single arrow.*

Proof. We have already shown that for such a topos, the representable functors $a\text{Hom}(-, C)$ form a strong generating set of supercompact objects. To prove the converse, observe that if \mathcal{E} has a generating set of supercompact objects \mathcal{G} , the topology on \mathcal{G} defined in Lemma 3.1 will be of regular type. \square

The regular objects in a topos will always be closed under pullbacks; the proof of this is similar to the corresponding argument for coherent objects (see Lemma 3.6).

Theorem 4.14. *A topos \mathcal{E} is locally regular if and only if it is equivalent to $\mathbf{Sh}(\mathcal{C}, J)$, where (\mathcal{C}, J) is a site for which \mathcal{C} has pullbacks and J is of regular type.*


Proof. It remains to show that in this situation, the functors $a\text{Hom}(-, C)$ will be regular for every object C in \mathcal{C} . The proof of this is similar to the proof of Lemma 3.9. \square

4.2 Modular lattices and m-dimension

Definition 4.15. *A lattice L is modular if for all $x, y, z \in L$ with $z \leq x$,*

$$x \wedge (y \vee z) = (x \wedge y) \vee z.$$

This is equivalent to the statement that L does not contain a pentagon [14, IV.1.1], that is, a sublattice of the form



In particular, a distributive lattice is modular, so any results we prove about modular lattices will also apply to distributive ones.

We quote the following result from [14, IV.1.2].

Theorem 4.16 (Diamond Isomorphism Theorem). *Let L be a modular lattice, and let $a, b \in L$. Then the mapping $x \mapsto x \wedge b$ defines an isomorphism between the intervals $[a, a \vee b]$ and $[a \wedge b, b]$ in the lattice.*

Recall that a congruence on a lattice L is an equivalence relation θ on L such that whenever $a \theta b$ in L , then for any $x \in L$, $(x \wedge a)\theta(x \wedge b)$ and $(x \vee a)\theta(x \vee b)$. This allows us to define the quotient lattice L/θ to be the set of equivalence classes, with (well-defined) lattice operations inherited from those on L .

For an interval $[a, b]$ in L , denote by $\text{con}(a, b)$ the congruence on L generated by (a, b) - this is the intersection of all congruences θ on L such that $a\theta b$.

Terminology surrounding the congruences on a lattice is described in [14, III.1], and can be used to describe the congruence generated by an interval $[a, b]$ (for $a, b \in L$, the *interval* $[a, b]$ is the set $\{x \in L \mid a \leq x \leq b\}$). Let L be a lattice. Suppose we have intervals $[a, b]$, $[c, d]$ in L , such that $b = a \vee d$ and $c = a \wedge d$. If this is the case, we say the intervals $[a, b]$ and $[c, d]$ are *perspective* to one another. The Diamond Isomorphism Theorem is the statement that in a modular lattice, intervals perspective to one another are isomorphic. If $[a, b]$ and $[c, d]$ are perspective to one

another, and θ is a congruence identifying a and b , then θ identifies c and d also. Perspectivity defines a reflexive, symmetric relation on the set of intervals in L . We call the transitive closure of this relation projectivity - that is, two intervals $[a, b]$ and $[c, d]$ are said to be *projective* to one another if there is a finite series of intervals $[c, d] = [x_0, y_0], [x_1, y_1], \dots, [x_n, y_n] = [a, b]$ such that each $[x_i, y_i]$ is projective to $[x_{i+1}, y_{i+1}]$. If two intervals $[a, b]$ and $[c, d]$ are projective then $\text{con}(a, b) = \text{con}(c, d)$.

For an interval $[a, b]$ in L , the congruence $\text{con}(a, b)$ identifies every pair of points in $[a, b]$. We expand our notion of projectivity to account for this. An interval $[c, d]$ is *weakly perspective into the interval* $[a, b]$ if there is some subinterval $[a', b']$ of $[a, b]$ such that $[c, d]$ is perspective to $[a', b']$. Weak projectivity is the transitive closure of this property: we say the interval $[c, d]$ is *weakly projective into* $[a, b]$ if there is a finite series of intervals $[c, d] = [x_0, y_0], [x_1, y_1], \dots, [x_n, y_n] = [a, b]$ such that each $[x_i, y_i]$ is weakly perspective into $[x_{i+1}, y_{i+1}]$. If $[c, d]$ is weakly projective into $[a, b]$, then any congruence θ identifying a and b must also identify c and d .

The notion of weak projectivity is enough to describe the congruence generated by an interval.

Theorem 4.17. (*[14, III.1.2]*) *Let L be a lattice, and let $[a, b], [c, d]$ be intervals in L . Then $(c, d) \in \text{con}(a, b)$ if and only if there is a sequence $c = x_0 \leq x_1 \leq \dots \leq x_n = d$ such that each interval $[x_i, x_{i+1}]$ is weakly projective into $[a, b]$.*

The development of m -dimension that follows is from [23, 7.1].

Let \mathcal{L} be a class of modular lattices closed under sublattices and quotients. For a given modular lattice L , let $\theta^0 = \theta_{\mathcal{L}}^0$ be the smallest congruence on L which collapses all intervals in L which occur in \mathcal{L} (we can consider an interval $[a, b]$ in L to be a modular lattice with top element b and bottom element a). Now set $L_0 = L$, and $L_1 = L_0/\theta^0$, with canonical map $\pi_1 : L_0 \rightarrow L_1$.

Continuing in this way, for each α , assuming L_α has been defined, we define θ^α to be the congruence on L_α generated by all the intervals in \mathcal{L} , and so on. Associated with this, we define a congruence θ_α on L to be the collection of intervals collapsed by the canonical map $L \rightarrow L_{\alpha+1}$. We use this to extend the definition to limit ordinals,

by defining $\theta_\lambda = \bigcup_{\alpha < \lambda} \theta_\alpha$.

Let α be the least ordinal (if there is one) such that L_α is the trivial lattice. We say $\alpha - 1$ is the \mathcal{L} -dimension of the lattice L (note that α cannot be a limit ordinal, for if, given a limit ordinal λ , the top and bottom elements of L are identified in θ_λ , they must be identified in some θ_α with $\alpha < \lambda$). If there is no such α , we say L has \mathcal{L} -dimension ∞ . The map $Q_\alpha : L \rightarrow L_\alpha$ is called the α -th \mathcal{L} -derivative.

We refer to a lattice with only two elements as a *gap*; such a lattice clearly has the structure $\circ \rightarrow \circ$.


Definition 4.18. *If \mathcal{L} consists of just gaps and trivial lattices (equivalently, all finite length lattices) we refer to the dimension it determines as m -dimension.*

For any lattice L , denote by $\mathbf{Cong}(L)$ the lattice of congruences on L ordered by inclusion; this is clearly a complete distributive lattice. We show that the congruences used to define m -dimension can be identified just by looking at the structure of $\mathbf{Cong}(L)$. An element x of a lattice L is called an *atom* or a *minimal element* if the interval $[0, x]$ in L is a gap.


Lemma 4.19. *Let L be a modular lattice with m -dimension $< \infty$, and let θ be a congruence relation on L . Then θ is an atom in $\mathbf{Cong}(L)$ if and only if it is a congruence generated by a gap in L .*

Proof. (\Leftarrow) Let $[a, b]$ be a gap in L , and let $x \in L$, such that $a \wedge x \neq b \wedge x$. Then $[a \wedge x, b \wedge x]$ is also a gap, by the diamond isomorphism theorem. For let $y = b \wedge x$; then $a \wedge y = a \wedge (b \wedge x) = a \wedge x$ and $b \vee y = b$.

That is, the diagram:



is the same as the diagram



and taking $x = y$, we can write any $y \in L$ with $a \vee y = b$ in the first form. Moreover, the interval $[a \wedge y, y]$ is perspective back to $[a, b]$ by taking $b \vee (-)$.

A similar argument holds when we take $[x \vee a, x \vee b]$.

So we have shown that any (non-trivial) interval onto which $[a, b]$ is weakly projective must be a gap, and moreover, each such gap is weakly projective back to $[a, b]$.

The result now follows from [14, Thm. 2, p.172]; any interval $[x, y]$ collapsed by the congruence $\text{con}(a, b)$ must contain a subinterval weakly projective to $[a, b]$; say $[x', y']$, $x \leq x' \leq y' \leq y$; $\text{con}(x', y')$ then collapses $[a, b]$ (since we can reverse the projection) and so we must have

$$\text{con}(a, b) \leq \text{con}(x', y') \leq \text{con}(x, y) \leq \text{con}(a, b)$$

(\Rightarrow) Let θ be a minimal congruence, and let $[a, b]$ be an interval collapsed by θ .

Suppose $[a, b]$ contains a gap, $[x, y]$ say. Then by minimality, the congruence generated by collapsing the interval $[x, y]$ is θ .

Therefore we must assume $[a, b]$ does not contain a gap. But if this is the case then for every $x, y \in [a, b]$ with $x \leq y$, there must be some $z \in [a, b]$ with $x < z < y$. Therefore $[a, b]$ is a densely ordered subset of L ; by [23, 7.2.3], the m -dimension of L is ∞ . □

This lemma tells us that to find the congruence θ^0 used in the construction of m -dimension, it is sufficient to take the join of all the atoms in $\mathbf{Cong}(L)$. Furthermore, note that congruences on $L_1 = L/\theta^0$ correspond to congruences on L containing θ^0 ; that is, they correspond to elements in the interval $[\theta_0, 1]$ in $\mathbf{Cong}(L)$. Thus, repeating the earlier result, we can find θ_1 - it is the join of all the atoms in the

lattice $[\theta_0, 1]$.

In general let L be a complete modular lattice. By the *socle* of L , we mean the join of all the atoms; we write $\mathbf{soc}(L)$. Define a series of elements of L : set $x_0 = \mathbf{soc}(L)$. Assuming that $x_\alpha \in L$ has been defined, set $x_{\alpha+1} = \mathbf{soc}([x_\alpha, 1])$. For limit ordinals λ , set $x_\lambda = \bigvee_{\alpha < \lambda} x_\alpha$. Let α be the least ordinal such that $x_\alpha = 1$, if such an ordinal exists. Then we say that the *socle length* of L is α . If no such ordinal exists, we say that the socle length of L is ∞ .

We have shown:

Corollary 4.20. *Let L be a modular lattice on which m -dimension is defined. Then the socle length of $\mathbf{Cong}(L)$ is equal to $\text{mdim}(L)$.*

Lemma 4.21. *Let L be a complete modular lattice; let $x \in L$. The socle length of the lattice $[x, 1]$ is bounded above by the socle length of L .*

Proof. Let $y_0 = 0$, $x_0 = x$. For each ordinal α , $y_{\alpha+1} = \mathbf{soc}[y_\alpha, 1]$, $x_{\alpha+1} = \mathbf{soc}[x_\alpha, 1]$, and for limit ordinals λ , $x_\lambda = \bigvee_{\alpha < \lambda} x_\alpha$, $y_\lambda = \bigvee_{\alpha < \lambda} y_\alpha$. It suffices to show that for each ordinal α , $y_\alpha \leq x_\alpha$.

Assume we've shown this for α . Let a be an atom in $[y_\alpha, 1]$. Then either $a \vee x_\alpha = x_\alpha$, or $a \vee x_\alpha$ is atom in $[x_\alpha, 1]$ (since the interval $[x_\alpha, a \vee x_\alpha]$ is isomorphic to $[y_\alpha, a]$ in L). It follows that $x_\alpha \vee \mathbf{soc}[y_\alpha, 1] \leq \mathbf{soc}[x_\alpha, 1]$ in L , since for each atom $a \in [y_\alpha, 1]$, $x_\alpha \vee a \leq \mathbf{soc}[x_\alpha, 1]$. Thus $y_{\alpha+1} \leq x_{\alpha+1}$.

The case for limit ordinals is straightforward. □

4.3 Krull-Gabriel dimension for abelian categories

The notions of \mathcal{L} -dimension for lattices obtained by taking congruences generated by the intervals which are in \mathcal{L} can be extended to give us a notion of dimension for abelian categories. We note that for any object A in an abelian category, the subobject lattice $\text{Sub}(A)$ is modular.

Let \mathcal{C} be a locally coherent abelian category, and let \mathcal{A} be the full subcategory of

coherent (equivalently, in locally coherent abelian categories, finitely presented) objects in \mathcal{C} . Let \mathcal{L} be a class of modular lattices closed under quotients and subobjects. The Serre subcategory $\mathcal{S}_{\mathcal{L}}$ of \mathcal{A} is that generated by the objects of \mathcal{A} whose subobject lattices (in \mathcal{A}) are in \mathcal{L} . The corresponding torsion class will be denoted $\mathcal{T}_{\mathcal{L}}$.

Bearing in mind Theorem 2.28, the following lemma [23, 13.1.2] looks a bit like the compactness theorem for regular logic; it is a case of Lemma 2.22.

Lemma 4.22. *Let \mathcal{L} be a class of modular lattices closed under sublattices and quotient lattices. Let \mathcal{A} be an abelian category and let $F \in \mathcal{A}$. Then $F \in \mathcal{S}_{\mathcal{L}}$ if and only if F has a finite chain of finitely generated subfunctors*

$$0 = F_0 < F_1 < \dots < F_n = F$$

such that for each i , the lattice $\text{Sub}(F_{i+1}/F_i)$ belongs to \mathcal{L} .

Now fix a locally coherent category \mathcal{C} ; write \mathcal{A} for the subcategory of coherent objects in \mathcal{C} , as above. Fix a class of lattices \mathcal{L} .

Write \mathcal{S}_0 for the Serre subcategory of \mathcal{A} generated by those objects subobject lattices are in \mathcal{L} . This gives rise to a quotient functor $Q_1 : \mathcal{A} \rightarrow \mathcal{A}_1$, which in turn corresponds to a finite type localization $\mathcal{C} \rightarrow \mathcal{C}_1$.

In general, assume we have a quotient functor $Q_\alpha : \mathcal{A} \rightarrow \mathcal{A}_\alpha$. Let \mathcal{S}^α be the Serre subcategory of \mathcal{C}_α determined by \mathcal{L} ; this gives us a localization $Q_{\alpha+1} : \mathcal{A} \rightarrow \mathcal{A}_{\alpha+1} = \mathcal{A}_\alpha/\mathcal{S}_{\mathcal{L}}$. Let $\mathcal{S}_{\alpha+1}$ be the kernel of $Q_{\alpha+1}$; this is a Serre subcategory of \mathcal{A} . For limit ordinals λ , take $\mathcal{S}_\lambda = \bigcup_{\alpha < \lambda} \mathcal{S}_\alpha$; this gives us a localization $Q_\lambda : \mathcal{A} \rightarrow \mathcal{A}_\lambda$.

This gives us a system of quotient functors $Q^\alpha : \mathcal{A} \rightarrow \mathcal{A}_\alpha$, or equivalently, a system of finite type localizations $Q_\alpha : \mathcal{C} \rightarrow \mathcal{C}_\alpha$. We call the quotient functor Q_α the α th \mathcal{L} -derivative.

Definition 4.23. *Let C be an object of \mathcal{C} ; then the \mathcal{L} -dimension of C is the least ordinal α such that $Q_{\alpha+1}(C) = 0$, if such an ordinal exists, and ∞ otherwise.*

Note that for a limit ordinal λ , if $Q_\lambda(C) = 0$, then $C \in \mathcal{T}_\lambda$, and consequently $C \in \mathcal{T}_\alpha$ for some $\alpha < \lambda$, so λ cannot be the least ordinal with $Q_\lambda(C) = 0$.

Definition 4.24. *The \mathcal{L} -dimension of \mathcal{C} is the least ordinal α such that $\mathcal{C}_{\alpha+1}$ is a groupoid (i.e., equivalent to the one-object category), if such an α exists, and ∞ otherwise.*

The following result [23, 13.1.4] explains the connection between the \mathcal{L} -dimensions for abelian categories and the lattice dimensions mentioned earlier.

Theorem 4.25. *Let \mathcal{L} be a class of modular lattices closed under sublattices and quotient lattices, let \mathcal{C} be a locally coherent abelian category with subcategory \mathcal{A} of coherent objects, and let A be an object of \mathcal{A} . Then $\mathcal{L}\text{-dim}(A) = \mathcal{L}\text{-dim}(\text{Sub}_{\mathcal{A}}(A))$.*

In the particular case where \mathcal{L} consists just of the one- and two-element lattices, we refer to the dimension we get as Krull-Gabriel dimension.

Definition 4.26. *The Krull-Gabriel dimension of a locally coherent abelian category \mathcal{C} is the \mathcal{L} -dimension where \mathcal{L} consists solely of the one and two object lattices.*

If \mathcal{L} consists of the one and two object lattices, we refer to the \mathcal{L} -derivatives on \mathcal{C} as the *Krull-Gabriel derivatives*.

We call an object A in \mathcal{A} *simple* if it is non-zero and has no non-trivial finitely presented subobjects. The next result is another case of Lemma 2.22.

Lemma 4.27. *Let A be a simple object of \mathcal{A} ; then the Serre subcategory \mathcal{S}_A generated by A consists precisely of objects F in \mathcal{A} such that there is a finite composition series*

$$0 = F_0 < F_1 \dots < F_n = F$$

such that for each $i \leq n$, the quotient F_i/F_{i-1} is isomorphic to A .

Corollary 4.28. *Let \mathcal{C} be a locally coherent abelian category for which Krull-Gabriel dimension is defined, and let \mathcal{A} be the subcategory of coherent objects in \mathcal{C} . A Serre subcategory of \mathcal{A} is an atom in the lattice of Serre subcategories of \mathcal{A} if and only if \mathcal{A} is generated by a simple object of \mathcal{A} .*

Proof. It is clear that the Serre subcategory generated by a simple object is minimal.

To see that a minimal Serre subcategory must be generated by a simple object, let \mathcal{S} be a minimal Serre subcategory. If \mathcal{S} contains a simple object A , it must be generated by it - for it contains the Serre subcategory generated by A and is minimal. So assume \mathcal{S} contains no simple objects.

Let $F \in \mathcal{S}$. Let S and S' be subobjects of F such that $S < S'$. Since $F \in \mathcal{S}$, we have that $S, S' \in \mathcal{S}$, and therefore $S/S' \in \mathcal{S}$. Since S/S' is not simple, it must have a proper subobject, which corresponds to an object X lying properly between S and S' in $\text{Sub}(F)$. Consequently, the lattice $\text{Sub}(F)$ is a dense ordered set, and so does not have m -dimension. But by Theorem 4.25, this means F does not have Krull-Gabriel dimension, and therefore \mathcal{A} does not either. \square

Corollary 4.29. *For a locally coherent abelian category \mathcal{C} , the Krull-Gabriel dimension of \mathcal{C} is equal to the socle length of the lattice of finite type localizations of \mathcal{C} .*

4.4 Krull-Gabriel dimension for toposes

In this section we try to provide a definition of Krull-Gabriel dimension for toposes. We would like a definition which for classifying toposes, corresponds to some associated property of the underlying theory, much as the correspondence between Krull-Gabriel dimension and m -dimension holds in the abelian case. Furthermore, for additive theories, the Krull-Gabriel dimension of the associated classifying topos ought to be the same as the Krull-Gabriel dimension of the associated additive functor category, since both correspond to properties of the underlying theory.

In the abelian case, we took the Serre subcategory of $\mathbf{fp}(\mathcal{C})$ generated by the simple objects in that category. By Theorem 2.27, we know that when \mathcal{C} is the functor category $(R\text{-Mod}, \mathbf{Ab})$ for some ring R , these simple objects are the pp-pairs ϕ/ψ , where $[\psi, \phi]$ is a simple interval in the lattice pp_R . Since the coequaliser diagram

$$\psi \rightarrow \phi \rightarrow \phi/\psi$$

is preserved by the localization, the inclusion $\psi \rightarrow \phi$ is mapped to an isomorphism

by it. Moreover, it is clear that any localization which inverts this map also contains ϕ/ψ in its kernel.

The inclusion $\psi \rightarrow \phi$ here has the property that it does not factor through any other monomorphisms - call such a monomorphism *simple*. From the previous paragraph, we see the following result:

Lemma 4.30. *Let $Q_0 : \mathcal{C} \rightarrow \mathcal{C}_1$ be the first Krull-Gabriel derivative for an abelian category \mathcal{C} . Then Q_0 is the smallest localization which inverts all the simple monomorphisms in $\mathbf{fp}(\mathcal{C})$.*

From this it seems reasonable to propose a definition for toposes as follows: the Krull-Gabriel dimension of the topos \mathcal{E} is to be the dimension associated with the derivative, which at each stage α is the Grothendieck topology generated by all the simple monomorphisms in a certain subcategory, such as the category of finitely presented, coherent or regular objects in \mathcal{E} .

Unfortunately, in the topos case, there is a problem with all of the above ideas. Consider a theory of algebras with an identity elements, e.g., groups, or R -modules for some ring R . Then in the classifying topos $(\mathbf{fp}(\mathbb{T}\text{-mod}), \mathbf{Sets})$, the initial object $\mathbf{0}$ is regular (because it always is) and the terminal object $\mathbf{1}$ is also (it is represented by the formula $x = 0$). The inclusion $\mathbf{0} \rightarrow \mathbf{1}$ is simple - let F be any subobject of $\mathbf{1}$, and suppose A is some algebra such that FA is non-empty. For any other algebra B there is a zero map $0 : A \rightarrow B$, and considering $F0(FA)$, we see that FB is non-empty as well. Therefore FB is non-empty for every algebra B , and so $F \cong \mathbf{1}$.

Now let F be any object in the classifying topos. The diagram below is a pullback:

$$\begin{array}{ccc} \mathbf{0} & \longrightarrow & F \\ \downarrow & & \downarrow \\ \mathbf{0} & \longrightarrow & \mathbf{1} \end{array}$$

Any localization of \mathcal{C} which inverted the map $\mathbf{0} \rightarrow \mathbf{1}$ would have to invert every map $\mathbf{0} \rightarrow F$ as well. Thus, any definition of Krull-Gabriel dimension similar to that discussed previously would give all theories of modules Krull-Gabriel dimension 0.

Corollary 4.29 suggests an alternative approach - we define Krull-Gabriel dimension for a topos to be the socle length of the lattice of coherent type localizations of the topos (i.e., for which the inclusion functor preserves directed colimits). From the Duality Theorem, if \mathcal{E} is the classifying topos for a coherent theory \mathbb{T} , these will be precisely the localizations that correspond to coherent quotients of \mathbb{T} . In fact, this too has a problem with it.

Consider the theory of sets. By a coherently definable subcategory of this category we mean one which is a model of some coherent theory. This is equivalent to being closed under ultraproducts and pure subobjects [27, 3.3]. In particular, directed colimits can be expressed as ultraproducts, so coherently definable categories are closed under directed colimits. It can be easily seen that every monomorphism $A \rightarrow B$ in **Sets** with $A \neq \mathbf{0}$ is pure, and every monic $\mathbf{0} \rightarrow B$ with $B \neq \mathbf{0}$ is not.

The whole category **Sets** is obviously coherently definable within itself, as is the empty subcategory (specified by the axiom \perp). For each natural number n the category \mathcal{S}_n consisting of all sets of size less than n , is coherently definable, and specified by the axiom

$$\forall x_1 \dots x_{n+1} \bigvee_{1 \leq i \neq j \leq n+1} (x_i = x_j).$$

Adding the axiom $\exists x(x = x)$ specifies the subcategory $\mathcal{S}_n \setminus \{\mathbf{0}\}$. Considering this last axiom in isolation gives us the coherently definable subcategory $\mathbf{Sets} \setminus \{\mathbf{0}\}$. Finally, the subcategory $\{\mathbf{0}\}$ containing just the empty set is a definable subcategory, specified by the axiom $\exists x(x = x) \rightarrow \perp$.

We claim these are all the coherently definable categories of **Sets**. Since such categories are closed under pure monics, it is clear that such categories are closed under non-zero subobjects, so it is enough to look at what the largest set in such a category is. Assume we have a coherently definable subcategory \mathcal{D} , and assume \mathcal{D} contains a countable set D . Let λ be any cardinal. Then the set $X = \coprod_{i \leq \lambda} D_i$, where each D_i is the size of D , may be expressed as the directed colimits of all the finite coproducts it contains. The finite coproducts are all countable, and therefore

isomorphic to D , and so are contained in \mathcal{D} ; thus so too is X . It follows that \mathcal{D} contains sets of arbitrary size.

The lattice of coherently definable subcategories of **Sets** therefore looks like this:

$$\begin{array}{ccccccc}
 \emptyset & \longrightarrow & \mathcal{S}_1/\{\mathbf{0}\} & \longrightarrow & \dots & \longrightarrow & \mathcal{S}_n/\{\mathbf{0}\} & \longrightarrow & \dots & \longrightarrow & \mathbf{Sets}/\{\mathbf{0}\} \\
 \downarrow & & \downarrow & & & & \downarrow & & & & \downarrow \\
 \{\mathbf{0}\} & \longrightarrow & \mathcal{S}_1 & \longrightarrow & \dots & \longrightarrow & \mathcal{S}_n & \longrightarrow & \dots & \longrightarrow & \mathbf{Sets}
 \end{array}$$

The lattice of coherent type localizations of the classifying topos of sets is the opposite of this lattice. It's clear that the socle of this is the localization corresponding to the subcategory $\mathbf{Sets}/\{\mathbf{0}\}$. However the lattice we get in the second step has no minimal quotient - that is, for the theory of sets, the lattice of coherent type localizations of the classifying topos has no socle length.

This example shows that there is a problem with trying to define Krull-Gabriel dimension using coherent type localizations. The problem is essentially that while regular formulas can allow us to discuss whether a given model of a theory has elements with certain properties, coherent formulas allow us to talk about how many elements it might have with this property. To avoid this, we can instead consider the regular type localizations. Recall that the localizations of a topos can be ordered by inclusion of the corresponding Grothendieck topologies with some particular site. Equivalently, if $a : \mathcal{E} \rightarrow \mathcal{F}$ and $a' : \mathcal{E} \rightarrow \mathcal{F}$ are two localizations of a topos \mathcal{E} , then $a \leq a'$ if the functor a' factors through a . We make the following definition:

Definition 4.31. *Let \mathcal{E} be a regular topos (i.e., one which is the classifying topos for a regular theory). The Krull-Gabriel dimension of \mathcal{E} is the socle length of the lattice of regular type localizations of \mathcal{E} .*

By the Duality Theorem (Theorem 2.31), if \mathcal{E} is the classifying topos for a regular theory \mathbb{T} , this definition corresponds to the dimension defined by the socle length of the lattice of regular quotients of \mathbb{T} .

Lemma 4.32. *If \mathcal{E} is a locally supercompact topos, and $a : \mathcal{E} \rightarrow \mathcal{F}$ is a regular type localization, then $KGdim(\mathcal{F}) \leq KGdim(\mathcal{E})$.*

Proof. This is immediate from Lemma 4.21. \square

We can see straight away that for an additive theory, for example the theory of modules over a ring, the Krull-Gabriel dimension of the classifying topos will equal the Krull-Gabriel dimension of the ring, since both are the socle length of the opposite of the lattice of regularly definable subcategories of $R\text{-Mod}$.

We also see from the previous example that the Krull-Gabriel dimension of the classifying topos of the theory of sets is 0 since, of the coherently definable subcategories we described there, the only ones which were regularly definable were \mathcal{S}_1 , $\mathcal{S}_1/\{\mathbf{0}\}$, **Sets** and **Sets**/ $\{\mathbf{0}\}$. This fits with our expectation that the theory of sets ought to have a very simple dimension.

4.5 The Ziegler frame

Let R be a ring. A regularly definable subcategory of $R\text{-Mod}$ is one which can be axiomatized by regular sequents in the language of R -modules (that is, it can be defined using only pp-formulas). The Ziegler spectrum of R is a topological space whose closed sets correspond to the regularly definable subcategories of $R\text{-Mod}$. For a construction of the Ziegler spectrum for a ring, see [23, ch.5].

We would like to generalise the construction of the Ziegler spectrum for a ring to non-additive theories. In this section, we look at a way to approach this problem.

Let Σ be a many-sorted signature of algebras. The regular theories over Σ form a complete lattice. Let $\{\mathbb{T}_i\}_{i \in I}$ be a collection of regular theories over Σ , that is, a collection of sets of regular sequents over Σ , closed under deduction. The join of these theories $\bigvee_{i \in I} \mathbb{T}_i$ is the deductive closure of the union $\bigcup_{i \in I} \mathbb{T}_i$. The meet $\bigwedge_{i \in I} \mathbb{T}_i$ is the set-wise intersection; this is easily seen to be deductively closed.

Lemma 4.33. *Let Σ be a many-sorted signature of algebras. The regular theories over Σ form a complete Heyting algebra.*

Proof. We have already observed that the regular theories over Σ form a complete lattice. So let \mathbb{T} , \mathbb{T}_i be regular theories over Σ , for i ranging over some indexing set

I. It suffices to verify the distributive law

$$\mathbb{T} \wedge \left(\bigvee_{i \in I} \mathbb{T}_i \right) = \bigvee_{i \in I} (\mathbb{T} \wedge \mathbb{T}_i).$$

Let σ be a sequent in $\bigvee_{i \in I} (\mathbb{T} \wedge \mathbb{T}_i)$. Then there is a collection of sequents $\sigma_1, \dots, \sigma_n$ admitting a deduction of σ such that for each $j = 1, \dots, n$ there is some $i_j \in I$ with $\sigma_j \in \mathbb{T} \wedge \mathbb{T}_{i_j}$. It follows that each σ_j is in $\bigvee_{i \in I} \mathbb{T}_i$, and so σ is in $\bigvee_{i \in I} \mathbb{T}_i$, since it is deductively closed. Similarly, σ is in \mathbb{T} , so $\sigma \in \mathbb{T} \wedge (\bigvee_{i \in I} \mathbb{T}_i)$.

To prove the converse, suppose σ is a sequent in $\mathbb{T} \wedge (\bigvee_{i \in I} \mathbb{T}_i)$. Then there is a finite collection of sequents $\sigma_j \in \mathbb{T}_{i_j}$, $1 \leq j \leq m \in \mathbb{N}$, admitting a deduction of σ .

Suppose σ has the form $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$, and each σ_j has the form $\phi_j(\mathbf{x}_j) \vdash \psi_j(\mathbf{x}_j)$. Write $\sigma_j \vee \sigma$ for the sequent $\phi_j \wedge \phi(\mathbf{x}_j, \mathbf{x}) \vdash \psi_j \vee \psi(\mathbf{x}_j, \mathbf{x})$.

We note that for each j , $\sigma_j \vee \sigma$ is in the theory $\mathbb{T} \vee \mathbb{T}_{i_j}$. This is true since there is a deduction of $\sigma_j \vee \sigma$ from both σ_j and σ .

Now consider the deduction of σ from the collection of sequents $\{\sigma_j\}_{j=1, \dots, m}$. We claim we can use this to write a deduction of σ from the collection of sequents $\sigma_j \vee \sigma$.

For each rule of deduction for regular logic, we claim we can replace any of the assumptions τ with $\tau \vee \sigma$, and write down a deduction of $\tau' \vee \sigma$, where τ' was the consequence in the original deduction.

As an example, suppose we have an instance of the cut rule in the deduction of σ , that is, a deduction:

$$\frac{\chi(\mathbf{y}) \vdash \eta(\mathbf{y}) \quad \eta(\mathbf{y}) \vdash \theta(\mathbf{y})}{\chi(\mathbf{y}) \vdash \theta(\mathbf{y})}.$$

If we replace the sequent $\chi(\mathbf{y}) \vdash \eta(\mathbf{y})$ with $\chi \wedge \phi(\mathbf{y}, \mathbf{x}) \vdash \eta \vee \psi(\mathbf{y}, \mathbf{x})$, we can write down the deduction

$$\frac{\chi \wedge \phi(\mathbf{y}, \mathbf{x}) \vdash \eta \vee \psi(\mathbf{y}, \mathbf{x}) \quad \frac{\eta(\mathbf{y}) \vdash \chi(\mathbf{y})}{\eta \vee \psi(\mathbf{y}, \mathbf{x}) \vdash \chi \vee \psi(\mathbf{y}, \mathbf{x})}}{\chi \wedge \phi(\mathbf{y}, \mathbf{x}) \vdash \eta \vee \psi(\mathbf{y}, \mathbf{x})}$$

One can easily verify that every instance of a rule of deduction in the derivation of σ from τ_1, \dots, τ_n can be replaced by a corresponding deduction to obtain a derivation

of $\sigma \vee \sigma$ from $\tau_1 \vee \sigma, \dots, \tau_n \vee \sigma$; alternatively, we can appeal to the Completeness Theorem for Regular Logic to argue that there must be such a deduction.

The deduction of σ from the sequents $\sigma_j \vee \sigma$ proves that σ is in $\bigvee_{i \in I} (\mathbb{T} \wedge \mathbb{T}_i)$. \square

This argument can be generalized to show that the distributive law holds for all geometric theories over the signature; that the geometric theories over signature form a Heyting algebra is shown explicitly in [10, 5.3]. Indeed, we could alternatively have proved this by noting that the Heyting operator constructed in [10, 5.7] will still define a Heyting operator if restricted to the regular sequents.

Definition 4.34. *Let \mathbb{T} be a regular theory over a many-sorted signature of algebras Σ . The Ziegler frame $\mathcal{Z}_{\mathbb{T}}$ is the lattice of regular quotients of the theory \mathbb{T} .*

The next result is the equivalent of saying that the space associated with the Ziegler frame has a basis of compact open sets.

Lemma 4.35. *Let \mathbb{T} be a regular theory over a many-sorted signature of algebras; then $\mathcal{Z}_{\mathbb{T}}$ has a generating set of finitely generated objects.*

Proof. The quotient theories \mathbb{T}' adding a single sequent $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$ to \mathbb{T} form a generating set for $\mathcal{Z}_{\mathbb{T}}$. This is true since any theory is simply the union of all the single-sequent theories it contains.

A quotient of \mathbb{T} defined by adding a single sequent σ of the form $\phi(\mathbf{x}) \vdash \psi(\mathbf{x})$ is finitely generated as an object in this lattice, because any deduction of σ from a collection of other quotient theories \mathbb{T}'_i uses only finitely many sequents, and thus only a finite number of the theories \mathbb{T}'_i are required. \square

For module categories, it can be shown that the pp-definable subcategories are determined by the pure-injective indecomposable modules they contain [23, 5.1.4]

It is always possible to choose a set of objects \mathcal{Z} in an algebraic category such that the regularly definable subcategories are determined by their intersection with \mathcal{Z} - given any pair of distinct regularly definable subcategories, we can find an object that is in one but not the other. It is an open problem whether this can be done in any ‘nice’ sort of way, as is the case with modules over a ring. In the next section,

we find such a collection of objects for the category of G -sets, where G is an abelian group. We will use this to calculate the Krull-Gabriel dimension for the classifying topos of the theory of G -sets.

4.6 A Ziegler spectrum for abelian groups

For the rest of this chapter, we look at how to compute the Krull-Gabriel dimension defined previously, for the classifying topos of the theory of G -sets where G is some abelian group. Such a category is the most obvious non-additive analogue of a category of modules. To help with the computations, we construct a ‘Ziegler spectrum’ - a set of objects in the category sufficient to specify regular theories, in the sense that a regularly definable subcategory is determined by its intersection with the spectrum. If such a set \mathcal{Z} exists, the intersections of \mathcal{Z} with the regularly definable subcategories will form the closed sets of a topology, justifying the use of the term spectrum.

We will use the following notation: since the group G is abelian, we will use additive notation for its product, and if X is a G -set, we will write $x + g$ for the image of an element $x \in X$ under the action of an element $g \in G$. For a set of elements \mathbf{g} from G , we define

$$\text{fix}_X(\mathbf{g}) = \{x \in X \mid \forall g \in \mathbf{g}(x + g = x)\}.$$

The paper [28] contained a construction of such a spectrum; unfortunately, there is a mistake in Lemma 4.3 of this paper. This result stated that the coherently definable subcategories of a category of G -sets are characterized by the objects they contain which have no pairs of isomorphic subobjects.

In the proof in [28] it is stated that if X and X' are G -sets for a group G , such that for every subset S of G , $\text{fix}_X(S) \neq \emptyset \leftrightarrow \text{fix}_{X'}(S) \neq \emptyset$, then X and X' satisfy precisely the same coherent sequents. This is not the case. To see an explicit counterexample, let G be the trivial group, so the category of G -sets is just the category **Sets**. A coherently definable subcategory of **Sets** is just the category of sets containing less

than n elements, for some $n \in \mathbb{N}$ (and of course, the category **Sets** itself). However, since there are no proper subgroups of the trivial group, any non-empty set X would satisfy the above condition.

The approach described in the paper [28] cannot be made to classify coherently definable subcategories, for this reason. It can however be used to construct a spectrum which classifies regularly definable subcategories. We describe this construction in the rest of this section.

Let G be an abelian group. G encodes an algebraic theory in one sort, consisting of unary functions corresponding to the elements of the group, and equations stating that composition of functions agrees with the group multiplication. The models of this theory are sets with a G -action on the left; we write G -**sets** for the category of such models. The classifying topos for this theory is $(\mathbf{fp} \ G\text{-sets}, \mathbf{Sets})$, the category of set-valued functors on the finitely presented G -sets.

We wish to examine the regularly definable subcategories in G -**sets**. Regularly definable subcategories are those closed under products, directed colimits and pure subobjects. To study them we want to understand the structure of products and pure embeddings in this category.

Every G -set X can be expressed as a disjoint union of orbits. We call a G -set *indecomposable* if it consists only of a single orbit. Such a set has the form G/H where H is some subgroup of G .

Whenever we have two indecomposable G -sets G/H and G/H' , where H and H' are subgroups of G , then each element $(x, y) \in G/H \times G/H'$ has as its stabilizer $H \cap H'$, so $G/H \times G/H'$ is a disjoint union of copies of $G/H \cap H'$. Similarly we see that given an arbitrary collection $\{G/H_i\}_{i \in I}$ of indecomposable G -sets, their product $\prod_{i \in I} G/H_i$ is a disjoint union of copies of $G/(\bigcap_{i \in I} H_i)$.

Now taking products and coproducts in G -**sets** preserves the underlying set structure, so given G -sets A, B, X and Y , the elements of $(A \amalg B) \times (X \amalg Y)$ are precisely the pairs (a, x) with a in A or B , and x in X or Y ; that is, we have the distributive

law

$$(A \coprod B) \times (X \coprod Y) = (A \times X) \coprod (A \times Y) \coprod (B \times X) \coprod (B \times Y).$$

In particular, if all these are indecomposable G -sets, we have that $(G/H_1 \coprod G/H_2) \times (G/H_3 \coprod G/H_4)$ has indecomposable components given in the following table:

	G/H_3	G/H_4
G/H_1	Copies of $G/(H_1 \cap H_3)$	Copies of $G/(H_1 \cap H_4)$
G/H_2	Copies of $G/(H_2 \cap H_3)$	Copies of $G/(H_2 \cap H_4)$

We can see from this how the indecomposable components of a product of G -sets $A \times B$ relate to the indecomposable components of A and B .

Pure monomorphisms of G -sets are characterized in [28, 3.1].

Lemma 4.36. *An embedding $X \rightarrow Y$ in G -sets is pure if and only if for every finite tuple \mathbf{g} from G , $\text{fix}_Y(\mathbf{g}) \neq \emptyset$ implies $\text{fix}_X(\mathbf{g}) \neq \emptyset$*

In particular, if every indecomposable component occurring in Y occurs in X , then the inclusion $X \rightarrow Y$ is pure.

Let X be an indecomposable G -set; thus X consists of only a single orbit, and for any $x, x' \in X$, there is some $g \in G$ such that $x + g = x'$. Now let $f : Y \rightarrow X$ be any map of G -sets, with Y non-empty; then there is some x in X that is in the image of Y , say $x = f(y)$ for $y \in Y$. For any $x' = x + g \in X$, we have that $f(y + g) = x'$. Thus f is a surjection. If f is also an injection, it must be an isomorphism (since toposes are balanced). Thus there are no non-trivial monomorphisms between indecomposable G -sets, and we see that every monomorphism of G -sets has the form $X \rightarrow X \coprod Y$.

We are now in a position to specify a set of objects sufficient to determine the regularly definable subcategories of G -sets. However we will not use this as our definition of a Ziegler spectrum, because as we shall see, we can in fact improve this result further to requiring only a subset of this set. Let $\{H_i : i \in I\}$ be the set of all proper subgroups of G up to conjugacy, and set $H_1 = H_2 = G$ (we assume $1, 2 \notin I$). Let \mathcal{Y} be the collection of G -sets

$$\mathcal{Y} = \left\{ \coprod_{i \in I'} G/H_i : I' \subseteq I \cup \{1, 2\} \right\}.$$

Proposition 4.37. *Let G be an abelian group. Every regularly definable subcategory ξ of G -sets is uniquely determined by its intersection with \mathcal{Y} .*

Proof. Let ξ be a regularly definable subcategory, and let X be an object in ξ . Let X' be the largest subobject of X contained in \mathcal{Y} .

Each indecomposable component in X occurs in X' . To see this, let $X = \coprod_{j \in J} X_j$, where each X_j is an indecomposable summand. These indecomposable summands can be collected into isomorphism classes, and X' is simply a coproduct of one representative from each of these classes, with an extra orbit isomorphic to $\mathbf{1}$ if X has more than one such orbit. The inclusion of X' into X is pure, and so X' is in ξ .

We now show that any regularly definable subcategory ξ' containing X' must also contain X . Set $Y = (X')^{|J|}$. Let $X_j \neq \mathbf{1}$ be an indecomposable summand in X (and so in X'); then Y contains a copy of $X_j^{|J|}$, and this G -set in turn contains (at least) $|J|$ copies of X_j . If X' contains $\mathbf{2}$ as a direct summand, then we can repeat the above argument to show that Y contains at least one copy of $\mathbf{1}$ for each indecomposable summand in X . If X' contains $\mathbf{1}$ but not $\mathbf{2}$, then Y contains precisely one copy of $\mathbf{1}$, but so too does X . Thus, there is an inclusion $X \rightarrow Y$; we have shown that Y has enough orbits of each type to guarantee that we can choose such an inclusion.

We claim that this inclusion is pure. Let \mathbf{g} be a finite tuple from G , and let $x = \{x_k\}_{k \in J} \in X'^{|J|} = Y$ be such that $gx = x$ for all g from \mathbf{g} . The stabilizer of x is the intersection of the stabilizers of the x_k , that is, $\text{Stab}(x) = \bigcap_{k \in J} \text{Stab}(x_k)$. So for any particular choice of j , we have that $\mathbf{g} \subseteq \text{Stab}(x) \subseteq \text{Stab}(x_k)$, so $x_k \in X' \subseteq X$ is also a fixed point for \mathbf{g} . This shows the inclusion is pure.

Now let ξ' be any regularly definable subcategory containing X' . Then ξ' must contain Y , since it is closed under products. Since X is a pure subobject of Y , ξ' must contain X also. So any regularly definable subcategory containing X' must also contain X ; the result is an immediate consequence of this. \square

Remark. The assertion that G is abelian is necessary here. For instance, take G to

be S_3 , given by the presentation $\{a, b \mid a^2b = ba, a^3 = b^2 = 1\}$. The G -set $G/\langle b \rangle$ has three elements, given by $a\langle b \rangle$, $ba\langle b \rangle$ and $b^2a\langle b \rangle$, whose stabilizers are $\langle a \rangle$, $\langle ba \rangle$ and $\langle b^2a \rangle$ respectively. Since each of these stabilizers has precisely one fixed point, we can write a regular sequent stating that a given G -set X contains at most one copy of this orbit:

$$(a(x) = a(x')) \vdash (x = x')$$

Consequently, regularly definable subcategories cannot be determined by their intersection with \mathcal{Y} in this case, because this intersection cannot tell us whether objects are allowed to contain one copy of this orbit, or more than one.

An object X in a category \mathcal{C} is *pure-injective* if whenever there is a pure embedding $m : A \rightarrow B$ in \mathcal{C} and a morphism $f : A \rightarrow X$, there is a (not necessarily unique) map $g : B \rightarrow X$ with $gm = f$.

$$\begin{array}{ccc} A & \xrightarrow{m} & B \\ f \downarrow & \swarrow g & \\ X & & \end{array}$$

We write $\text{pinj}(\mathcal{C})$ for the collection of all pure-injective objects in \mathcal{C} .

A map $m : X \rightarrow Y$ is *pure-essential* if it is pure and whenever $f : Y \rightarrow Z$ is a map such that fm is a pure morphism, then f is monic. A G -set Y is the *pure-injective hull* of X if it is pure-injective and there is a pure-essential map $m : X \rightarrow Y$. This is equivalent to stating that any pure morphism $f : X \rightarrow Y'$ with Y' pure-injective factors through m via a monomorphism $Y \rightarrow Y'$.

The pure-injective hull is unique up to isomorphism. For suppose $m : X \rightarrow Y$, $m' : X \rightarrow Y'$ are both pure-injective hulls. Then there are maps $f : Y \rightarrow Y'$, $f' : Y' \rightarrow Y$ such that $fm = m'$ and $f'm' = m$; the maps f and f' are both monic by the pure-essential property.

The existence of the pure-injective hull in the category $G\text{-sets}$ is shown in [5]; this construction was used in [28] also. We describe the construction now. We will follow the notation of [28].

For a G -set X , let $FF(X)$ be the set of all subgroups H of G such that for every finitely generated subgroup $H' \leq H$, $\text{fix}_X(H') \neq \emptyset$. Let \mathbb{M}_X be the subset of $FF(X)$ consisting of those subgroups H that are maximal in $FF(X)$, and have no fixed points in X . Let $I\mathbb{M}_X$ be a set of representatives of conjugacy classes of \mathbb{M}_X .

Proposition 4.38. ([5, 3.4]) *The G -set $X \coprod (\coprod_{H \in I\mathbb{M}_X} G/H)$ is a pure-injective hull for X .*

The pure-injective hull of X is elementarily equivalent to X [28, 3.6]. Note that if X is in \mathcal{Y} then its pure-injective hull $H(X)$ is in \mathcal{Y} too, since we adjoin only a collection of orbits which are isomorphic neither to each other nor to any orbits already in X . Furthermore, if ξ is a regularly definable subcategory and X is in ξ , then $H(X)$ is in ξ too, since it is elementarily equivalent to X . So a regularly definable subcategory is determined by the pure-injectives it contains that are in \mathcal{Y} . We have now proved:

Lemma 4.39. *Let G be an abelian group. Let ξ be a regularly definable subcategory of G -sets. Then ξ is determined by its intersection with $\mathcal{Z} = \mathcal{Y} \cap \text{pinj}(G\text{-sets})$.*

This allows us to make the following definition:

Definition 4.40. *Let G be an abelian group. The Ziegler spectrum of the theory of G -sets is given by $\mathcal{Z}_G = \mathcal{Y}_G \cap \text{pinj}(G\text{-sets})$ (that is, the collection of pure-injective G -sets with no two isomorphic indecomposable components, save for allowing two copies of the one-element G -set $\mathbf{1}$).*

4.7 An example with a cyclic group

Let \mathbb{Z}_n be a cyclic group of finite order. Every subgroup of \mathbb{Z}_n is of the form \mathbb{Z}_m , where $m \mid n$. Suppose for simplicity that n has no square divisors, that is, n may be written as the product of distinct primes.

As in the previous section, the elements of \mathcal{Y} can be partially ordered by pure inclusion. In particular, $Y \leq Y'$ in \mathcal{Y} if and only if for every indecomposable component G/H in Y , there is an isomorphic indecomposable component in Y' , and for

An indecomposable \mathbb{Z}_n -set necessarily has the form \mathbb{Z}_m . Define the *degree* of \mathbb{Z}_m to be the number of distinct prime factors of m . For a general \mathbb{Z}_n -set X , let the degree of X be the largest number d such that X has an indecomposable component of degree d .

The following lemma allows us to determine which subsets of \mathcal{Y} correspond to maximal regular subcategories.

Lemma 4.41. *If an object X in \mathcal{Y} has degree $m > 1$ then it can be expressed as a pure subobject of a product of objects in \mathcal{Y} of lower degree.*

Proof. Suppose X has only one indecomposable component of degree d , of the form $\mathbb{Z}_{p_1 \dots p_d}$, i.e., $X = X' \amalg \mathbb{Z}_{p_1 \dots p_d}$, where X' has degree $< d$. Then X embeds purely into $\prod_{i=1}^d (X' \amalg \mathbb{Z}_{p_1 \dots \hat{p}_i \dots p_d})$. For suppose $x = (x_i)_{i=1}^d$ is an element of this product. Either $x_i \in X'$ for some i , in which case $\text{Stab}(x) \subseteq \text{Stab}(x_i)$, or $x_i \in \mathbb{Z}_{p_1 \dots \hat{p}_i \dots p_d}$ for every i , in which case $\text{Stab}(x) = \frac{n}{p_1 \dots p_d}$.

In general, assume X has r such components. Write X as the coproduct

$$X = X' \amalg \left(\prod_{j=1}^r \mathbb{Z}_{p_{j_1} \dots p_{j_d}} \right)$$

where X' has degree $< d$. Then consider the G -set

$$\tilde{X} = \prod_{1 \leq k_1 \leq i_1, \dots, 1 \leq k_r \leq i_r} \left(X' \amalg \left(\prod_{j=1}^r \mathbb{Z}_{p_{j_1} \dots \hat{p}_{j_{k_j}} \dots p_{j_d}} \right) \right)$$

There is an embedding $X \rightarrow \tilde{X}$, and this embedding will be pure, using a similar argument to the case where there is only one indecomposable component of degree d . □

For each set of distinct primes p_1, \dots, p_n , consider the axiom

$$\exists x_1, \dots, x_n \left(\bigwedge_{i=1}^n x_i + p_i = x_i \right) \vdash \exists y \left(y + \frac{n}{p_1 \dots p_n} = y \right)$$

This defines a subcategory $\zeta(\mathbb{Z}_{p_1} \amalg \dots \amalg \mathbb{Z}_{p_n})$ of \mathbb{Z}_n -Sets whose intersection with \mathcal{Y} consists of all the objects in \mathcal{Y} save those in the lattice $[\mathbb{Z}_{p_1} \amalg \dots \amalg \mathbb{Z}_{p_n}]$. This is

a maximal regularly definable subcategory, because any strictly larger subcategory must contain as objects $\mathbb{Z}_{p_1} \amalg \dots \amalg \mathbb{Z}_{p_n}$ and all the elements in $[\mathbf{1}]$; by taking products we see that such a category must contain the whole of \mathcal{Y} .

Thus every object in \mathcal{Y} of the form $\mathbb{Z}_{p_1} \amalg \dots \amalg \mathbb{Z}_{p_n}$ must be excluded from the intersection of all the maximal regularly definable subcategories.

Any object X of \mathcal{Y} that does not contain a copy of $\mathbf{1}$ not of the above form must be expressible as a (pure subobject of) a product of objects of \mathcal{Y} which are of this form. So let $X \subseteq \prod_{i \in I} X_i$ be such an expression. Let ζ be any maximal regularly definable subcategory. If any of the X_i 's do not appear in ζ , then we must have $\zeta = \zeta(X_i)$, the regularly definable subcategory whose intersection with \mathcal{Y} is the whole of \mathcal{Y} save the lattice $[X_i]$. In particular, this implies $X \in \zeta$. If all of the X_i 's are in ζ , then X must also be in ζ , since ζ is closed under products.

To summarise the above: a maximal regularly definable subcategory is specified by taking all but one of the lattices in \mathcal{Y} . The excluded lattice is not allowed to be $[\mathbf{1}]$.

Further maximal definable subcategories can be found by taking non-trivial intersections with $[\mathbf{1}]$. For each prime p dividing n , consider the axiom:

$$(x + p = x) \wedge (z + 1 = z) \vdash (x + \frac{n}{p} = x).$$

This defines a subcategory ζ , containing all the elements of \mathcal{Y} not in $[\mathbf{1}]$ (since if there is no indecomposable component isomorphic to $\mathbf{1}$, no witness for z can be found), and whose intersection with the lattice $[\mathbf{1}]$ consists of precisely those G -sets not containing an isomorphic copy of \mathbb{Z}_p . This ζ is a maximal regularly definable subcategory - to see this, suppose we have some G -set in \mathcal{Y}/ζ , say $\mathbf{1} \amalg \mathbb{Z}_p \amalg X$. Clearly $\mathbf{1} \amalg \mathbb{Z}_p$ embeds purely into this. Now let $\mathbf{1} \amalg \mathbb{Z}_p \amalg Y$ be any other \mathbb{Z}_n -set in \mathcal{Y}/ζ . Then Y contains no indecomposable component isomorphic to \mathbb{Z}_p , so $\mathbf{1} \amalg Y$ is in ζ , and moreover there is a pure embedding of $\mathbf{1} \amalg \mathbb{Z}_p \amalg Y$ into $(\mathbf{1} \amalg Y) \times (\mathbf{1} \amalg \mathbb{Z}_p)$. Thus any regularly definable subcategory containing both ζ and $\mathbf{1} \amalg \mathbb{Z}_p \amalg X$ contains the whole of \mathcal{Y} , and therefore is the whole of the category \mathbb{Z}_n -sets.

Therefore the intersection of all the maximal regularly definable subcategories, χ_1 , is given by the objects in \mathcal{Y} which don't contain any indecomposable component of prime order.

Now let X be any object of $\chi_1 \cap \mathcal{Y}$ in which the order of every indecomposable component has at most two prime factors, e.g., $X = \mathbb{Z}_{m_1} \amalg \dots \amalg \mathbb{Z}_{m_n}$, with each m_i having at most two prime factors.

Then consider the axiom

$$\exists x_1, \dots, x_n \left(\bigwedge_{i=1}^n x_i + m_i = x_i \right) \vdash \exists y \left(y + \frac{n}{\text{lcm}(m_1 \dots m_n)} = y \right)$$

This combined with the axioms of χ_1 gives a regularly definable subcategory of \mathbb{Z}_n -Sets whose intersection with \mathcal{Y} is all the objects in $\chi_1 \cap \mathcal{Y}$ save those appearing in the lattice $[X]$.

Likewise, for each pair of prime factors p_1, p_2 of n , the axiom

$$(x + p_1 p_2 = x) \wedge (z + 1 = z) \vdash \left(x + \frac{n}{p_1 p_2} = x \right)$$

defines a subcategory of \mathbb{Z}_n -sets whose intersection with $\chi_1 \cup \mathcal{Y}$ contains all the objects except the elements of the lattice $[1]$ containing an indecomposable component isomorphic to $\mathbb{Z}_{p_1 p_2}$.

Thus we can show that χ_2 , the intersection of all the maximal regularly definable subcategories of \mathbb{Z}_n -Sets contained in χ_1 , is the subcategory generated by all those objects in \mathcal{Y} containing an indecomposable component whose order has at least three distinct prime factors.

Continuing in this way, we eventually show:

Proposition 4.42. *Let $n \in \mathbb{N}$ be such that n has no square divisors. The Krull-Gabriel dimension of the classifying topos for the theory of \mathbb{Z}_n -sets is equal to the number of distinct prime factors of n .*

Chapter 5

Flat Covers in Algebraic Categories

In this chapter, we look at the flat cover conjecture, which claims that every object in a category has a flat cover. This result was first proposed for categories of modules in [13], and eventually proven by Bican, El Bashir and Enochs in [6]. This result has been proven for more general additive categories in, for example, [31], [12], [19]. The approach used for additive categories has to a certain extent been generalized in [30]. We shall use these ideas to study the notion of a flat cover in a topos. Using the notion of a coessential epimorphism as the definition for a cover, Renshaw and Mahmoud studied flat covers in categories of monoid acts, in [21]. In this paper they asked when such a category has covers in the previous sense; this is the question we address in this chapter.

5.1 Strongly flat models of a theory

Given an S -sorted equational theory (Σ, E) , we can define an associated category \mathcal{T} , which is closed under finite products. The objects of \mathcal{T} are the finite products of sorts from S . The morphisms

$$X_1 \times \dots \times X_n \rightarrow X'_1 \times \dots \times X'_m$$

are precisely the m -tuples of terms $(t_1(\mathbf{x}), \dots, t_m(\mathbf{x}))$ over n -tuples of variables $\mathbf{x} \in X_1 \times \dots \times X_n$, where for each $j = 1, \dots, m$, $t_j(\mathbf{x})$ is a term of sort X'_j . We stipulate that for each equation $t(\mathbf{x}) = t'(\mathbf{x})$ in E , the corresponding morphisms in \mathcal{T} are equal (as are morphisms containing these terms as components, and equal in the other components).

A model of the equational theory (Σ, E) then becomes the same thing as a functor $\mathcal{T} \rightarrow \mathbf{Sets}$ preserving finite products.

This gives rise to a more general way of looking at algebraic theories, which is studied in detail in [3].

Definition 5.1. *An algebraic theory is a small category \mathcal{T} with finite products.*

An algebra for the theory \mathcal{T} is a functor $A : \mathcal{T} \rightarrow \mathbf{Sets}$, preserving finite products.

A homomorphism of algebras is a natural transformation between functors.

The category of algebras associated with a theory \mathcal{T} is denoted $\mathbf{Alg} \mathcal{T}$.

A category \mathcal{C} is called algebraic if it is equivalent to $\mathbf{Alg} \mathcal{T}$ for some algebraic theory \mathcal{T} .

Lemma 5.2. ([3, 3.18]) *Algebraic categories are exact.*

In particular, this means that extremal, strong and regular epimorphisms all coincide in algebraic categories [8, 2.1.4].

Associated with algebraic theories, we have the following concept.

Definition 5.3. *A category \mathcal{D} is called sifted if finite products in \mathbf{Sets} commute with colimits over \mathcal{D} .*

A sifted colimit is a colimit over a sifted diagram.

An object C in a category \mathcal{C} is called perfectly presentable if $\text{Hom}(C, -)$ preserves sifted colimits.

Directed colimits are sifted [1, 1.59].

Perfectly presentable objects play a similar role in the study of algebraic categories as finitely presentable objects in the study of locally finitely presentable categories.

The next definition describes an important property of perfectly presentable objects in an algebraic category.

Definition 5.4. *An object P in a category \mathcal{C} is regular projective if it is projective with respect to regular epimorphisms; that is, given a diagram*

$$\begin{array}{ccc} & & P \\ & & \downarrow f \\ E & \xrightarrow{e} & A \end{array}$$

where e is a regular epimorphism, there is a map $f' : P \rightarrow E$ with $ef' = f$.

Lemma 5.5. ([3, 5.16]) *In an algebraic category, an object is perfectly presentable if and only if it is finitely presentable and regular projective.*

This allows us to write down a characterization of algebraic categories similar to that for locally finitely presentable categories.

Theorem 5.6. ([3, 18.4]) *A category \mathcal{C} is algebraic if and only if it is exact, co-complete and has a strong generating set \mathcal{G} of regular projective finitely presentable objects.*

Corollary 5.7. *Every algebraic category is locally finitely presented.*

In Lemma 2.1 we characterized finitely generated objects in a locally finitely generated category as strong quotients of a coproduct of objects from the generating set. Lemma 2.2 proved a similar condition characterizing finitely presented objects in a locally finitely presentable category. The next result provides an analogous characterization of perfectly presentable objects in an algebraic category. It is essentially the same result as [3, 5.14], which characterizes the perfectly presentable objects as split subobjects of representable algebras. We do not make clear the connection between choosing generating sets for the category and choosing an underlying algebraic theory for it here; this connection is detailed in the proof of [3, 6.9].

Lemma 5.8. *Let \mathcal{C} be an algebraic category, with a generating set \mathcal{G} of perfectly presentable objects. Then an object C in \mathcal{C} is perfectly presentable if and only if there is a split epimorphism*

$$\coprod_{i=1}^n G_i \begin{array}{c} \xrightarrow{e} \\ \xleftarrow{m} \end{array} C$$

where each object G_i is in \mathcal{G} .

Proof. (\Rightarrow) If such a diagram exists, the epimorphism e is regular, since it is split, and the kernel pair of e contains the morphisms $me, 1_{\coprod_{i=1}^n G_i}$. Thus C is a finite colimit of finitely presented objects, so it is finitely presented.

Since it is a coproduct of regular projective objects, $\coprod_{i=1}^n G_i$ is regular projective, and since C is a split subobject of a regular projective object it is regular projective also.

(\Leftarrow) If \mathcal{C} is perfectly presentable, there is an extremal epimorphism $e : \coprod_{i=1}^n G_i \rightarrow C$, which is regular, since in an exact category with pushouts, regular and extremal epimorphisms coincide ([8, 2.1.4], [1, 0.5]). Since C is regular projective, e admits a splitting. \square

Algebraic categories are a fairly general setting in which we can develop the idea of a strongly flat object.

Definition 5.9. *Let \mathcal{C} be any category. An epimorphism $e : E \rightarrow C$ in \mathcal{C} is pure if every map $f : A \rightarrow C$ with A finitely presented admits a factorization through e , that is, there is some map $g : A \rightarrow E$ with $eg = f$.*

Lemma 5.10. *Let \mathcal{C} be an algebraic category, and let C be an object in \mathcal{C} . Then the following are equivalent:*

1. C is a directed colimit of perfectly presentable objects.
2. There is a pure regular epimorphism $e : F \rightarrow C$ with F a coproduct of perfectly presentable objects.
3. Every morphism $A \rightarrow C$ with A finitely presented factors through some perfectly presentable object.
4. Every regular epimorphism $e : E \rightarrow C$ in \mathcal{C} is pure.

Proof. (1) \Rightarrow (3): if C is a directed colimit of perfectly presentable objects, then every morphism $G \rightarrow C$ with G finitely presented must factor through one of the objects in the directed system.

(3) \Rightarrow (1): since the category \mathcal{C} is locally finitely presented, C is expressible as a directed colimit of finitely presented objects. The factorization property (3) tells us that the perfectly presentable objects are cofinal in this system, and so C can be expressed as a directed colimit of these objects, see [1, 0.11].

(2) \Rightarrow (3): suppose F is given by a coproduct $\coprod_{i \in I} F_i$ with each F_i perfectly presentable. Then F is the directed colimit of all the finite coproducts of the F_i , and each finite coproduct is perfectly presentable. A morphism $f : G \rightarrow C$ with G finitely presentable must factor through e by purity, and so through one of the finite coproducts over the F_i by the finitely presentable property.

(1) \Rightarrow (2): since the perfectly presentable objects strongly generate the category \mathcal{C} , there is a strong epimorphism $e : \coprod_{h:G \rightarrow C} G_h \rightarrow C$, where for each $h : G \rightarrow C$ the object G_h is isomorphic to G and the component of e is given by h . Since e is a strong epimorphism and the category \mathcal{C} is exact, e is regular. Moreover, if C is the directed colimit of perfectly presentable objects, then every morphism $a : A \rightarrow C$ with A finitely presentable factors through some map $G \rightarrow C$ with $G \in \mathcal{G}$, and hence through e .

(2) \Rightarrow (4): suppose there is a pure epimorphism $f : F \rightarrow C$ with F a coproduct of perfectly presentable objects, and let $e : E \rightarrow C$ be any regular epimorphism. Then since F is regular projective, there's a map $f' : F \rightarrow E$ with $ef' = f$. Given any map $a : A \rightarrow C$ with A finitely presented, there is a factorization $a' : G \rightarrow F$ with $fa' = a$, since f is pure. The map $f'a' : G \rightarrow E$ then gives us the required factorization to show that e is also pure.

(4) \Rightarrow (2): we have already seen that there is a regular epimorphism $F \rightarrow C$ with F a coproduct of perfectly presentable objects; this is pure by assumption. \square

Remark. If \mathcal{G} is a particular strong generating set of perfectly presentable objects, then all the perfectly presentable objects mentioned in the statement of proof of Lemma 5.10 can be assumed to be finite coproducts of objects from \mathcal{G} .

Definition 5.11. *Let \mathcal{C} be an algebraic category. An object in \mathcal{C} with any of the above properties is called strongly flat.*

This terminology is based on that used for acts over a monoid.

Remark. For an arbitrary locally finitely presentable category \mathcal{C} , it would appear reasonable to choose any set of finitely presented objects \mathcal{G} , and attempt to use these objects to define a notion of strongly flat object in \mathcal{C} (i.e., in each of the above conditions, we replace ‘perfectly presentable object’ with ‘object from \mathcal{G} ’). We can in fact prove the equivalences of definitions (1)-(3). However, exactly which objects are ‘strongly flat’ according to this definition depends on what generating set \mathcal{G} is chosen - for instance, it is clear from condition (1) that if \mathcal{G} is chosen to be all finitely presentable objects, then every object in \mathcal{C} is strongly flat according to this definition.

Theorem 5.12. *The strongly flat objects in an algebraic category are closed under directed colimits, pure subobjects and arbitrary coproducts.*

Proof. Let D be an object in \mathcal{C} which is a directed colimit of strongly flat objects, and $a : A \rightarrow D$ is any map with A finitely presentable. The map a must factor through some strongly flat object in the directed system, and so must factor through some perfectly presentable object.

To see that strongly flat objects are closed under pure subobjects, suppose $s : S \rightarrow C$ is a pure monomorphism with C strongly flat. Let $a : A \rightarrow S$ be a morphism with A finitely presented. Then sa factors through some perfectly presentable object G , and we have a diagram

$$\begin{array}{ccc} S & \xrightarrow{s} & C \\ \uparrow a & & \uparrow g \\ A & \xrightarrow{a'} & G \end{array}$$

Now since A and G are both finitely presented, and s is pure, there is a map $g' : G \rightarrow S$ with $g'a' = a$. Thus, S satisfies condition (3).

Finally, we show that an arbitrary coproduct of strongly flat objects is strongly flat. It suffices to prove this for finite coproducts; the result for infinite coproducts will then follow by considering an infinite coproduct to be the directed colimit of all the finite coproducts over the same collection of objects.

Suppose F and F' are strongly flat objects in an algebraic category \mathcal{C} . Then there are directed diagrams $D_F : (I, \leq) \rightarrow \mathcal{C}$, $D_{F'} : (J, \leq) \rightarrow \mathcal{C}$, for directed partially

ordered sets (I, \leq) and (J, \leq) , where for every $i \in I$ and $j \in J$, $D_F i$ and $D_{F'} j$ are perfectly presentable, and these diagrams have colimit F and F' respectively. Now consider the set $I \times J$ with the partial order defined by $(i, j) \leq (i', j')$ if and only if $i \leq i'$ and $j \leq j'$. This partial order is directed: the join of a pair of elements (i, j) and (i', j') is simply $(i \vee i', j \vee j')$. We define a functor $D : (I \times J, \leq) \rightarrow \mathcal{C}$ by $D(i, j) = D_F i \amalg D_{F'} j$; if $(i, j) \leq (i', j')$, the map $D_{(i,j),(i',j')} : D_F i \amalg D_{F'} j \rightarrow D_F i' \amalg D_{F'} j'$ is given by $(D_F)_{ii'} \amalg (D_{F'})_{jj'}$.

Each of the objects $D_F i \amalg D_{F'} j$ is a finite coproduct of perfectly presentable objects, and so is itself perfectly presentable. The colimit of the diagram D is thus a directed colimit of perfectly presentable objects, and so is strongly flat. This colimit is easily seen to be isomorphic to $F \amalg F'$, and so $F \amalg F'$ is strongly flat. \square

Definition 5.13. ([3, 9.4, 9.7]) *A functor $G : \mathcal{A} \rightarrow \mathcal{B}$ between two algebraic categories is algebraic if it has a left adjoint and preserves sifted colimits.*

Theorem 5.14. *Let $G : \mathcal{A} \rightarrow \mathcal{B}$ be an algebraic functor between algebraic categories. Then the left adjoint $F : \mathcal{B} \rightarrow \mathcal{A}$ preserves perfectly presentable and strongly flat objects.*

Proof. Since F is a left adjoint, it preserves colimits. Since strongly flat objects are directed colimits of perfectly presentable objects, it suffices to show that F preserves perfectly presentable objects.

Let P be a perfectly presentable object in \mathcal{B} , and let $D : \mathcal{D} \rightarrow \mathcal{A}$ be a sifted diagram in \mathcal{A} . Write $\text{colim}(D)$ for the colimit of this diagram. Then we have natural isomorphisms

$$\begin{aligned} \text{Hom}_{\mathcal{A}}(FP, \text{colim}(D)) &\cong \text{Hom}_{\mathcal{B}}(P, G\text{colim}(D)) \cong \text{Hom}_{\mathcal{B}}(P, \text{colim}(GD)) \\ &\cong \text{colimHom}(P, -).GD \cong \text{colimHom}(FP, -).D \end{aligned}$$

This shows that the functor $\text{Hom}(FP, -)$ preserves sifted colimits, so FP is perfectly presented. \square

5.2 The Flat Cover Conjecture

Let \mathcal{C} be any category, and let \mathcal{F} be a subcategory.

Definition 5.15. • Let C be an object of \mathcal{C} . An \mathcal{F} -precover of C is a map $f : F \rightarrow C$ with F in \mathcal{F} , such that for any morphism $f' : F' \rightarrow C$ with F' in \mathcal{F} , there is a factorization of f' through f .

- An F -cover of C is an F -precover $f : F \rightarrow C$ with the additional property that for any endomorphism $e : F \rightarrow F$ with $fe = f$, e is an isomorphism.

An F -cover is unique up to isomorphism.

If every object in \mathcal{C} has an \mathcal{F} -precover, we say \mathcal{F} is *weakly coreflective* in \mathcal{C} . If every object has an \mathcal{F} -cover, we say \mathcal{F} is *stably weakly coreflective*. This terminology is from [30].

It can easily be shown that if \mathcal{C} has coproducts, every weakly coreflective subcategory of \mathcal{C} is closed under them.

Theorem 5.16. [30, 2.5] Let \mathcal{C} be a locally finitely presentable category, and \mathcal{F} a weakly coreflective full subcategory of \mathcal{C} which is closed under directed colimits in \mathcal{C} . Then \mathcal{F} is stably weakly coreflective in \mathcal{C} , i.e., every object in \mathcal{C} has an \mathcal{F} -cover.

Let \mathcal{C} be an algebraic category. We say \mathcal{C} satisfies the *Flat Cover Conjecture* if every object in \mathcal{C} admits a strongly flat cover; that is, if the subcategory \mathcal{SF} of strongly flat objects is stably weakly coreflective in \mathcal{C} .

It was shown in [6] that module categories satisfy the Flat Cover Conjecture; this result is generalized to locally finitely presentable abelian categories in [31, 5.2], and also in [12].

Since algebraic categories are locally finitely presented, and strongly flat objects in them are closed under directed colimits, it suffices to show that strongly flat precovers exist in order to prove the Flat Cover Conjecture for a given algebraic category.

5.3 Flat covers in presheaf categories

In this section we consider the existence of strongly flat covers for categories of set-valued functors over a small category. This provides the simplest case of a topos in which the existence of strongly flat covers can be investigated. One particular case of this will occur when the category in question has only one object, i.e., it is a monoid.

Let M be a monoid, with identity element 1_M . We may consider M to be a one object category, the functors $M \rightarrow \mathbf{Sets}$ are then precisely the *left M -acts*: an M -act is a set X with a multiplication $M \times X \rightarrow X$ satisfying $m.n(x) = m(n(x))$ and $1_M(x) = x$, for every $x \in X$ and every $m, n \in M$. Morphisms of M -acts are functions $f : X \rightarrow Y$ which preserve this structure, i.e., such that for every $m \in M$ and $x \in X$, the equation $f(m(x)) = m(f(x))$ holds; these are of course the natural transformations between the M -acts considered as functors. We denote the category of M -acts and their morphisms with the notation $M\text{-acts}$. This notation differs from that used earlier for G -sets, but its use is standard in semigroup theory. Strongly flat objects in categories of M -acts have already been investigated in [21], and in that paper it was asked what conditions on a monoid M can be used to show that the category $M\text{-acts}$ has strongly flat covers. At the end of this section, we will prove a condition for a category of set-valued functors to have strongly flat covers; this condition will include a significant class of monoids.

Let \mathcal{A} be a small category and let $(\mathcal{A}, \mathbf{Sets})$ be the category of set valued functors on \mathcal{A} . As described in section 3.1, we can associate with \mathcal{A} an essentially algebraic theory (in fact since all the operations will be total, we will get an equational presentation), whose models are the set-valued functors on \mathcal{A} . The finitely generated free objects for this theory will be the coproducts of representable functors. Since the category $(\mathcal{A}, \mathbf{Sets})$ is a topos, it is exact and all the epimorphisms are regular [20, IV.7.8]. The representable functors are projective (and therefore regular projective) and finitely presented, so they form a strong generating set of perfectly presented objects. Thus $(\mathcal{A}, \mathbf{Sets})$ is an algebraic category, and the representable functors form a strong generating set of perfectly presentable objects; in particular, by the remark

following Lemma 5.10, an object in $(\mathcal{A}, \mathbf{Sets})$ is strongly flat if and only if it is a directed colimit of finitely generated free objects.

Let F be a functor $\mathcal{A} \rightarrow \mathbf{Sets}$. Recall that the *category of elements* of F , $\int F$ has as objects pairs (x, A) where A is an object in \mathcal{A} and $x \in FA$, and as morphisms $(x, A) \rightarrow (x', A')$ morphisms $f : A \rightarrow A'$ in \mathcal{A} with $Ff(x) = x'$. The category of elements admits a forgetful functor $\pi_F : \int F \rightarrow \mathbf{Sets}$, given on objects by $(x, A) \mapsto FA$, and on morphisms by $[f : (x, A) \rightarrow (x', A')] \mapsto [Ff : FA \rightarrow FA']$.

Observe that taking the coproduct of functors is the same process as taking the disjoint union of the categories of elements - given a set of functors $\{F_i\}_{i \in I}$, we have that $\int \coprod_i F_i$ is equivalent to $\coprod_i \int F_i$. Furthermore, a natural transformation of functors $\alpha : F \rightarrow G$ gives a functor $\int \alpha : \int F \rightarrow \int G$, taking each object (x, A) in $\int F$ to $(\alpha_A(x), A)$ in $\int G$.

A small category \mathcal{C} is *connected* if the directed multigraph underlying \mathcal{C} is weakly connected. This is equivalent to saying that for any two objects of \mathcal{C} and \mathcal{C}' of \mathcal{C} , there is a sequence of objects $C = X_0, X_1, \dots, X_n = C'$ such that for each $i = 1, \dots, n$, either there is a morphism $f_i : X_{i-1} \rightarrow X_i$ or a morphism $f_i : X_i \rightarrow X_{i-1}$. The existence of such a sequence for a pair of objects defines an equivalence relation R on the objects of \mathcal{C} ; the equivalence classes of R are the connected components of \mathcal{C} . We can express \mathcal{C} as the coproduct of its connected components in the category \mathbf{Cat} of small categories.

Now let $F : \mathcal{A} \rightarrow \mathbf{Sets}$ be any functor. The connected components of $\int F$ correspond to certain subfunctors of F . For a given $(x, A) \in \int F$, define $F_x A' = \{y \in FA' \mid (y, A')R(x, A)\}$. This is a subfunctor of F , and $\int F_x$ is a connected component of $\int F$. Choosing a set S of representatives for the equivalence classes of R on $\int F$, we can write F as a coproduct $F = \coprod_{(x,A) \in S} F_x$, and each F_x will be indecomposable.

The image of an indecomposable functor under a natural transformation is indecomposable, since the image of the connected category \int_{F_x} under the corresponding functor must again be connected. In general, for a natural transformation $\alpha : F \rightarrow G$, we can write $F = \coprod_i F_i$ and $G = \coprod_j G_j$ with each F_i and G_j indecomposable, and for each i there is some j_i such that $\alpha(F_i) \subseteq G_{j_i}$.

A finitely generated functor F can have only finitely many components. For every functor is the directed colimit of its subfunctors which do have only finitely many components; if F is finitely generated it must be isomorphic to such a subfunctor.

Lemma 5.17. *An object F in $(\mathcal{A}, \mathbf{Sets})$ is strongly flat if and only if all of its indecomposable components are strongly flat.*

Proof. We have shown that strongly flat objects are closed under arbitrary coproducts in any category (Theorem 5.12).

Now suppose we have a collection of indecomposable objects F_i in $(\mathcal{A}, \mathbf{Sets})$, indexed by some set I , such that $F = \coprod_{i \in I} F_i$ is strongly flat. Suppose for some $i \in I$ we have a map $f : G \rightarrow F_i$ where G is finitely presented. Then the map $s_i f : G \rightarrow F$ obtained by composition with the inclusion $s_i : F_i \rightarrow F$ factors through some free object H . We have a diagram

$$\begin{array}{ccc} G & \xrightarrow{f'} & H \\ f \downarrow & & \downarrow h \\ F_i & \xrightarrow{s_i} & F \end{array}$$

The image of the map $h : H \rightarrow F$ is contained in F_i (since $(\mathcal{C}, \mathbf{Sets})$ is a topos, the inclusion maps into coproduct are monic, see e.g., [20, IV.10.5]), so we can write h as $s_i h'$ for some $h' : H \rightarrow F_i$. Now $s_i h' f' = s_i f$, so since s_i is monic, $h' f' = f$. This shows that f factors through H as $h' f'$, satisfying condition 3 for strong flatness. \square

Given a small category \mathcal{A} and a functor $F : \mathcal{A} \rightarrow \mathbf{Sets}$, we define a ‘Hom-functor’ $R_F : \mathbf{Sets} \rightarrow \mathbf{Sets}^{\mathcal{A}^{\text{op}}}$ by

$$(R_F X)(A) = \text{Hom}_{\mathbf{Sets}}(FA, X).$$

This functor has a left adjoint $L_F : \mathbf{Sets}^{\mathcal{A}^{\text{op}}} \rightarrow \mathbf{Sets}$ defined for each presheaf P by taking $L_F(P)$ to be the colimit of the diagram

$$\int P \xrightarrow{\pi_P} \mathcal{A} \xrightarrow{F} \mathbf{Sets}.$$

We will write $L_F(P)$ as $P \otimes F$. It is a set, whose elements are equivalence classes

of pairs $(x, y) \in PA \times FA$, where A is an object in \mathcal{A} , with the equivalence relation R generated by taking $(x, y)R(x', y')$ whenever there is map $f : A \rightarrow A'$ in \mathcal{A} with $Ff(y) = y'$ and $x'.Pf = x$.

Since it is a left adjoint, the tensor product commutes with colimits in $\mathbf{Sets}^{\mathcal{A}}$.

We define the following conditions for set-valued functors on \mathcal{A} , which are those used in the study of monoid acts to classify different levels of flatness.

Definition 5.18. *Let \mathcal{A} be a small category and suppose we have a functor $F : \mathcal{A} \rightarrow \mathbf{Sets}$.*

1. *F is said to satisfy condition (P) if given any $x \in FA$, $y \in FB$ and maps $f : A \rightarrow C$, $g : B \rightarrow C$ in \mathcal{A} with $Ff(x) = Fg(y)$, there is some $z \in D$ and maps $f' : D \rightarrow A$, $g' : D \rightarrow B$ in \mathcal{A} , such that $Ff'(z) = x$, $Fg'(z) = y$ and $ff' = gg'$.*
2. *F is said to satisfy condition (E) if given any $f, g : A \rightarrow B$ in \mathcal{A} and $x \in FA$ with $Ff(x) = Fg(x)$, there is some $h : C \rightarrow A$ in \mathcal{A} and some $y \in FC$ with $Fh(y) = x$ and $fh = gh$.*

These conditions allow us to write a further characterization of strong flatness for set-valued functors on a category. This is the characterization of strong flatness usually used by semigroup theorists writing about categories of monoid acts, see for example [18, III.9]. We verify here that the proof does work for arbitrary presheaf categories.

Lemma 5.19. *(cf. [32, 5.3], also [20, VII.6.3]) Let \mathcal{A} be a small category. For a functor $F : \mathcal{A} \rightarrow \mathbf{Sets}$, the following conditions are equivalent to F being a strongly flat object in the category $(\mathcal{A}, \mathbf{Sets})$.*

5. *F satisfies conditions (P) and (E).*
6. *The tensor product functor $(-) \otimes F$ preserves pullbacks and equalizers.*

Proof. We recall conditions (1) to (4) from Lemma 5.10.


(5) \Rightarrow (3): Suppose G is a finitely presentable functor, and there is a natural transformation $\alpha : G \rightarrow F$, where F satisfies conditions (P) and (E).

We want to show that the morphism α factors through some perfectly presentable object. It suffices to show this for the restriction of α to each irreducible component of G , since this will allow us to define a factorization for the whole of α using coproducts. Thus, we can assume that G is irreducible.

Suppose G has the presentation $\langle x_0, \dots, x_n \mid r_1, \dots, r_m \rangle$. Here each variable x_i is of sort A_i , and each relation r_j is of the form $f_1^j(x_{i_1}) = f_2^j(x_{i_2})$, where f_1 and f_2 are appropriate maps in \mathcal{A} (since every term in the theory corresponds to a morphism of \mathcal{A} , all relations in this theory must be of this form). Since G is irreducible, there is a reordering of the variables and relations such that for $i = 1, \dots, n$, the i th relation equates a term in the variable x_i and some variable x_k with $k < i$. If this were not the case, we could find two subsets \mathbf{x} and \mathbf{x}' of the variables such that every relation mentioned only terms in \mathbf{x} or \mathbf{x}' . We would then be able to represent G as the coproduct of the subobjects generated by \mathbf{x} and \mathbf{x}' respectively.

A map $G \rightarrow F$, corresponds to some tuple (y_1, \dots, y_n) in F satisfying all these relations. Now the relation r_1 is of the form $f_1^1(y_0) = f_2^1(y_1)$, so applying the (P) condition, there is some z_1 and arrow f'_1, f'_2 such that $f'_1(z_1) = y_0, f'_2(z_1) = y_1$, and $f'_1 f_1 = f'_2 f_2$. The relation r_2 involving y_2 and either y_0 or y_1 now implies a condition on z_1 and y_2 , say $f_1^2(z_1) = f_2^2(y_2)$. We apply the (P) condition again to obtain z_2 .

A picture at this point will help.



We can't fill in the arrows f_1^i since we don't know which of the y_i 's is going to be the source. After n steps, we get an element z_n which divides all of the original generators y_1, \dots, y_n . Now the relation r_{n+1} can be written as an equality

$f_1^{n+1}(z_{n+1}) = f_2^{n+1}(z_{n+1})$; applying condition (E) to this, we find z_{n+1} dividing z_n by h_{n+1} with $f_1^{n+1}h_{n+1} = f_2^{n+1}h_{n+1}$. Continuing, we get z_m dividing all the others, with z_m of sort A . Now the map $G \rightarrow F$ factors through the map $(A, -) \rightarrow F$ defined by $1_A \mapsto z_m$.

(6) \Rightarrow (5): Suppose we have $x \in FA, y \in FB$ and maps $f : A \rightarrow C, g : B \rightarrow C$ in \mathcal{A} with $Ff(x) = Fg(y)$. We construct a pullback diagram in $\mathbf{Sets}^{\mathcal{A}^{\text{op}}}$

$$\begin{array}{ccc} P & \xrightarrow{\pi_g} & (-, B) \\ \downarrow \pi_f & & \downarrow \cdot g \\ (-, A) & \xrightarrow{\cdot f} & (-, C) \end{array}$$

On an object D in \mathcal{A} , the set PD is the collection of all maps $D \rightarrow C$ admitting a factorization through both f and g . Tensoring with F gives a pullback diagram

$$\begin{array}{ccc} P \otimes F & \xrightarrow{\pi_f} & FB \\ \downarrow \pi_g & & \downarrow Fg \\ FA & \xrightarrow{Ff} & FC \end{array}$$

Since $Ff(x) = Fg(y)$, there exists some object D and $(h, z) \in PD \times FD$ such that h factorizes through f and g and the image of $z \in FD$ under these factorizations is x and y respectively; that is, z fulfils the condition (P).

The proof that condition (E) will be satisfied is similar.

(1) \Rightarrow (6): If F is a free functor, say $(-) \otimes (A, -)$, then any functor $P : \mathcal{A}^{\text{op}} \rightarrow \mathbf{Sets}$ is mapped to PA by this functor. In particular, since limits are calculated pointwise in $\mathbf{Sets}^{\mathcal{A}^{\text{op}}}$, this functor preserves all limits, and so in particular is strongly flat.

Now any strongly flat functor F can be constructed from free functors using finite coproducts and directed colimits. Since finite coproducts and directed colimits commute with equalizers and pullbacks in \mathbf{Sets} , it follows that $(-) \otimes F$ commutes with equalizers and pullbacks also. \square

Theorem 5.20. *Let \mathcal{C} be a category such that there is only a set of indecomposable strongly flat functors in $(\mathcal{C}, \mathbf{Sets})$. Then the category \mathcal{SF} of strongly flat functors is*

stably weakly coreflective in $(\mathcal{C}, \mathbf{Sets})$, i.e., $(\mathcal{C}, \mathbf{Sets})$ has \mathcal{SF} -covers.

Proof. By Lemma 5.12 and Theorem 5.16, it suffices to prove that \mathcal{SF} is weakly coreflective in M -acts.

Suppose $\{F_i \mid i \in I\}$ is a collection of representatives of all isomorphism classes of indecomposable strongly flat functors in $(\mathcal{C}, \mathbf{Sets})$.

Let X be a functor $\mathcal{C} \rightarrow \mathbf{Sets}$. Define $X^* = \coprod_{f \in (F_i, X), i \in I} (F_i)_f$, where each $(F_i)_f$ is isomorphic to F_i . Let $g : X^* \rightarrow X$ be the morphism defined on each $(F_i)_f$ by f . We claim that g is a precover of X . To see this, let $h : H \rightarrow X$ be any morphism, with H strongly flat. H is a coproduct of indecomposable strongly flat M -sets, say $H = \coprod_j H_j$. For each H_j , we have the map $H_j \longrightarrow H \xrightarrow{f} X$, which of course factors through X^* . The factorization of all of these maps through X^* gives a factorization of h through X^* . \square

To make any use of this result, we need to provide some examples of categories \mathcal{C} which have only a set of indecomposable strongly flat functors $F : \mathcal{C} \rightarrow \mathbf{Sets}$.


We acknowledge the help of Alex Bailey in formulating the next theorem.

Theorem 5.21. *Let \mathcal{C} be a category with the property that for any morphism $m : B \rightarrow C$ in \mathcal{C} , there is a natural number $k_m \in \mathbb{N}$ such that for any other morphism $s : A \rightarrow C$ there are at most k_m distinct morphisms $t : A \rightarrow B$ satisfying $mt = s$. The category $(\mathcal{C}, \mathbf{Sets})$ has strongly flat covers.*

Proof. It suffices to show that there is a regular cardinal λ such that for every indecomposable strongly flat functor F , the category $\int F$ has size less than λ (that is, $\coprod_{C \in \mathbf{ob}(\mathcal{C})} FC < \lambda$). If this is the case, then F considered as a model of the canonical theory over \mathcal{C} is λ -presented, and in particular there is only a set of such models up to isomorphism.

So let F be an indecomposable strongly flat functor, and let (x, C) be an object

in $\int X$. For any other object (y, B) in $\int X$ there is a path




Here each m_i is a morphism in \mathcal{C} .

Applying the (P) condition $\frac{n}{2} - 1$ times, we can reduce this setup to

$$\begin{array}{ccc}
 (x, A) & & (y, B) \\
 & \swarrow m & \nearrow m' \\
 & \circ &
 \end{array}
 \tag{5.1}$$

Thus for every (y, B) in $\int F$ there is some diagram of the above form. However, we will show that for each (x, A) in X and morphism $m : C \rightarrow A$ in \mathcal{C} , there are at most k_m elements $y \in FC$ with $Fm(z) = x$. To see this, suppose there are $k_m + 1$ distinct elements $y_1, \dots, y_{k_m+1} \in FC$ with $Fm(y_i) = x$ for each i . Applying the (P) condition to each pair in turn, we can find some object D in \mathcal{C} and some $z \in FD$ as shown



Now the $s_i \in \mathbf{mor}(\mathcal{C})$ are all distinct, and for each $i = 1, \dots, k_m + 1$, $ms_i = ms_1$, contradicting the condition that there are at most k_m distinct maps with this property.

Now given a choice of (x, A) in $\int F$, every other object in $\int F$ occurs in a diagram of the form (5.1), and there are only finitely many distinct occurrences of each such diagram. Thus there are at most $\aleph_0 |\mathbf{mor}(\mathcal{C})|^2$ objects in $\int F$. Between each pair of objects there are at most $|\mathbf{mor}(\mathcal{C})|$ morphisms, so $\int F$ has at most $\aleph_0 |\mathbf{mor}(\mathcal{C})|^3$ morphisms. Thus it suffices to take λ to be larger than $\aleph_0 |\mathbf{mor}(\mathcal{C})|^3$. \square

Examples of monoids with the property described above include *left cancellative monoids* (those monoids M such that for $s, s', t \in M$, $ts = ts'$ implies $s = s'$) and finite monoids.

A finitely generated monoid with the property described above is said to be of *(left) finite geometric type*. That is, a finitely generated monoid M is of left finite geometric type if for any $s \in M$ there is a natural number $k_s \in \mathbb{N}$ with the property that for any $t \in M$, there are at most k distinct elements $x_1, \dots, x_k \in M$ such that $sx_i = t$ for each $i = 1, \dots, k$. Monoids of finite geometric type have been studied in the literature of semigroups. We have shown that if a monoid M has finite geometric type, then the category M -**acts** satisfies the Flat Cover Conjecture.

Example. Let $M = (\mathbb{N}, +)$. To specify the action of $(\mathbb{N}, +)$ on a set, it suffices to specify the action of $1 \in \mathbb{N}$ on each element. Thus an $(\mathbb{N}, +)$ -act is specified by a directed graph in which each vertex is the source of precisely one arrow. A strongly flat $(\mathbb{N}, +)$ -act X cannot contain any loops when considered as a directed graph. A loop would correspond to an element $x \in X$ with some $n \in \mathbb{N} \setminus \{0\}$ such that $x + n = x$. This would contradict condition (E) - since $x + 0 = x + n$, we would be able to find some natural number m such that $m + n = m + 0 = m$. Similarly, an element of a strongly flat $(\mathbb{N}, +)$ act can have at most one divisor by 1. For suppose there exists y, y' with $y + 1 = x = y' + 1$. Then there is some $z \in X$ and elements $n, n' \in \mathbb{N}$ such that $z + n = y$, $z + n' = y'$, and $n + 1 = n' + 1$. But this would give us $n = n'$, and so $y = y'$. Thus the only indecomposable strongly flat $(\mathbb{N}, +)$ -acts are \mathbb{N} and \mathbb{Z} with the obvious actions of $(\mathbb{N}, +)$.

The map $\mathbb{Z} \rightarrow \mathbf{1}$ is an \mathcal{SF} -cover, where $\mathbf{1}$ is the one element $(\mathbb{N}, +)$ -act.

In the paper of Mahmoudi and Renshaw [21], a second definition of strongly flat cover was given. This defined a cover of an object C to be a *coessential epimorphism* $e : F \rightarrow C$, that is, an epimorphism e such that for any subobject $s : S \rightarrow F$, if es is an epimorphism then s is an isomorphism. Renshaw and Mahmoud asked whether the existence of a strongly flat cover in our sense implied the existence of such a coessential epimorphism from a strongly flat object. However, no such coessential epimorphism exists for $\mathbf{1}$ considered as an $(\mathbb{N}, +)$ -set; to see this, observe that any

map $\mathbb{N} \rightarrow \mathbf{1}$ must restrict to a subobject of \mathbb{N} on which it is also onto (or use [21, 3.5(1)]).

Bibliography

- [1] J. Adamek and J. Rosicky. *Locally Presentable and Accessible Categories*, volume 189 of *London Mathematical Society Lecture Notes*. Cambridge University Press, 1994.
- [2] J. Adamek and J. Rosicky. On pure quotients and pure subobjects. *Czechoslovak Mathematical Journal*, 58:623–636, 2004.
- [3] J. Adamek, J. Rosicky, and E. M. Vitale. *Algebraic Theories*, volume 184 of *Cambridge Tracts in Mathematics*. Cambridge University Press, 2011.
- [4] M. Artin, A. Grothendieck, and J-L. Verdier. *Séminaire de géométrie algébrique*, volume 269 of *Lecture Notes in Mathematics*. Springer, 1972.
- [5] B. Banachewski. Equational compactness of G-sets. *Canadian Mathematical Bulletin*, 17(1):11–18, 1974.
- [6] L. Bican, R. El Bashir, and E. Enochs. All modules have flat covers. *Bulletin of the London Mathematical Society*, 33:385–390, 2001.
- [7] F. Borceux. *Handbook of Categorical Algebra, vol. 1*, volume 50 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1994.
- [8] F. Borceux. *Handbook of Categorical Algebra, vol. 2*, volume 51 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1994.
- [9] K. Burke. *Some model-theoretic properties of functor categories for modules*. PhD thesis, University of Manchester, 1994.

- [10] O. Caramello. *Lattices of Theories*. Available on arXiv. http://arxiv.org/PS/_cache/arxiv/pdf/0905/0905.0299v1.pdf.
- [11] O. Caramello. *The Duality between Grothendieck Toposes and Geometric Theories*. PhD thesis, University of Cambridge, 2009.
- [12] S. Crivei, M. Prest, and B. Torrecillas. Covers in finitely accessible categories. *Proceedings of the American Mathematical Society*, 138:1213–1221, 2010.
- [13] E. Enochs. Injective and flat covers, resolvents and envelopes. *Israel Journal of Mathematics*, 39(3):189–209, 1981.
- [14] G. Grätzer. *General Lattice Theory (second edition)*. Birkhäuser, 2003.
- [15] R. Hartshorne. *Algebraic Geometry*. Springer, 1977.
- [16] P. Johnstone. *Sketches of an Elephant: A Topos Theory Compendium, vol. 1*. Oxford University Press, 2002.
- [17] P. Johnstone. *Sketches of an Elephant: A Topos Theory Compendium, vol. 2*. Oxford University Press, 2002.
- [18] M. Kilp, U. Knauer, and A. V. Mikhalev. *Monoids, Acts and Categories*, volume 29 of *Expositions in Mathematics*. de Gruyter, 1991.
- [19] H. Krause. Approximations and adjoints in homotopy categories. *Mathematische Annalen*. <http://dx.doi.org/10.1007/s00208-011-0703-y>.
- [20] S. MacLane and I. Moerdijk. *Sheaves in Geometry and Logic*. Springer, 1992.
- [21] M. Mahmoudi and J. Renshaw. On covers of cyclic acts over monoids. *Semigroup Forum*, 77(2):325–338, 2008.
- [22] N. Popescu. *Abelian Categories with Applications to Rings and Modules*. Academic Press, 1973.
- [23] M. Prest. *Purity, Spectra and Localisation*, volume 121 of *Encyclopedia of Mathematics*. Cambridge University Press, 2009.

- [24] M. Prest and A. Ralph. Locally finitely presented categories of sheaves of modules. *MIMS Eprints*, 2010. <http://eprints.ma.man.ac.uk/1410/>.
- [25] M. Prest and A. Ralph. On sheafification of modules. *MIMS Eprints*, 2010. <http://eprints.ma.man.ac.uk/1411/>.
- [26] R. Rajani. *Model-Theoretic Imaginaries and Ziegler Spectra in General Categories*. PhD thesis, University of Manchester, 2007.
- [27] R. Rajani and M. Prest. Model-theoretic imaginaries and coherent sheaves. *Applied Categorical Structures*, 17(6):517–559, 2009.
- [28] R. Rajani and M. Prest. Pure injectivity and model theory for G-sets. *Journal of Symbolic Logic*, 74(2):474–488, 2009.
- [29] J.-E. Roos. *Locally Noetherian categories and generalised strictly linearly compact rings. Applications*, volume 92 of *Lecture Notes in Mathematics*, pages 197–277. Springer-Verlag, 1969.
- [30] J. Rosicky. On projectivity in locally presentable categories. *Journal of Algebra*, (272):701–710, 2004.
- [31] W. Rump. Flat covers in abelian and non-abelian categories. *Advances in Mathematics*, 225:1589–1615, 2010.
- [32] B. Stenström. Flatness and localization over monoids. *Mathematische Nachrichten*, 48:315–334, 1971.