μ -ABSTRACT ELEMENTARY CLASSES AND OTHER GENERALIZATIONS

WILL BONEY, RAMI GROSSBERG, MICHAEL LIEBERMAN, JIŘÍ ROSICKÝ,
AND SEBASTIEN VASEY

ABSTRACT. We introduce μ -Abstract Elementary Classes (μ -AECs) as a broad framework for model theory that includes complete boolean algebras and metric spaces, and begin to develop their classification theory. Moreover, we note that μ -AECs correspond precisely to accessible categories in which all morphisms are monomorphisms, and begin the process of reconciling these divergent perspectives: for example, the preliminary classification-theoretic results for μ -AECs transfer directly to accessible categories with monomorphisms.

1. INTRODUCTION

In this paper, we offer a broad framework for model theory, μ -abstract elementary classes, and connect them with existing frameworks, namely abstract elementary classes and, from the realm of categorical model theory, accessible categories (see [MaPa89], [AdRo94]) and μ -concrete abstract elementary classes (see [LiRo]).

All of the above frameworks have developed in response to the need to analyze the model theory of nonelementary classes of mathematical structures; that is, classes in which either the structures themselves or the relevant embeddings between them cannot be adequately described in (finitary) first order logic. This project was well underway by the 50's and 60's, which saw fruitful investigations into infinitary logics and into logics with additional quantifiers (see [Dic75] and [BaFe85] for summaries). Indeed, Shelah [Sh702, p. 41] recounts that Keisler and Morley advised him in 1969 that this direction was the future of model theory and that

Date: February 11, 2016

AMS 2010 Subject Classification: Primary 03C48. Secondary: 03C45, 03C52, 03C55, 03C75, 03E55, 18C35.

Key words and phrases. Abstract elementary classes; Accessible categories; μ -Abstract elementary classes; Classification Theory; Tameness; Categoricity.

This material is based upon work done while the first author was supported by the National Science Foundation under Grant No. DMS-1402191.

The third and fourth authors are supported by the Grant Agency of the Czech Republic under the grant P201/12/G028.

This material is based upon work done while the fifth author was supported by the Swiss National Science Foundation under Grant No. 155136.

first-order had been mostly explored. The subsequent explosion in stability theory and its applications suggest otherwise, naturally, but the nonelementary context has nonetheless developed into an essential complement to the more classical picture.

On the model-theoretic side, Shelah was the leading figure, publishing work on excellent classes ([Sh87a] and [Sh87b]) and classes with expanded quantifiers [Sh43], and, of greatest interest here, shifting to a formula-free context through the introduction of *abstract elementary classes* (or *AECs*) in [Sh88]. The latter are a purely semantic axiomatic framework for abstract model theory that encompasses first order logic as well as infinitary logics incorporating additional quantifiers and infinite conjuncts and disjuncts, not to mention certain algebraically natural examples without an obvious syntactic presentation—see [BET07]. It is important to note, though, that AECs still lack the generality to encompass the logic L_{ω_1, ω_1} or complete metric structures.

Even these examples are captured by *accessible categories*, a parallel, but significantly more general, notion developed simultaneously among category theorists, first appearing in [MaPa89] and receiving comprehensive treatments both in [MaPa89] and [AdRo94]. An accessible category is, very roughly speaking, an abstract category (hence, in particular, not a category of structures in a fixed signature) that is closed under sufficiently directed colimits, and satisfies a kind of weak Löwenheim-Skolem property: any object in the category can be obtained as a highly directed colimit of objects of small size, the latter notion being purely diagrammatic and internal to the category in question. In particular, an accessible category may not be closed under arbitrary directed colimits, although these are almost indispensable in model-theoretic constructions: the additional assumption of closure under directed colimits was first made in [Ros97]—that paper also experimented with the weaker assumption of directed bounds, an idea that recurs in Section 6 below.

Subsequent work (see [BeRo12], [Lie11], and [LiRo]) has resulted in a precise characterization of AECs as concrete accessible categories with added structure, namely as pairs $(\mathcal{K}, U)$, where

- $\mathcal{K}$ is an accessible category with all morphisms monomorphisms and all directed colimits, and
- $U : \mathcal{K} \rightarrow \mathbf{Set}$ (with $\mathbf{Set}$ the category of sets and functions) is a faithful (“underlying set”) functor satisfying certain additional axioms.

Details can be found in Section 3 of [LiRo]. Of particular importance is the extent to which U preserves directed colimits; that is, the extent to which directed colimits are concrete. If we assume that U preserves arbitrary directed colimits, we obtain a category equivalent to an AEC. If we make the weaker assumption that U merely

preserves colimits of μ -directed, rather than directed, diagrams, we arrive at the notion of a μ -concrete AEC (see [LiRo]). Note that, although directed colimits may not be preserved by U (that is, they may not be “**Set**-like”), they still exist in the category $\mathcal{K}$ —metric AECs, whose structures are built over complete metric structures, are a crucial example of this phenomenon. One might ask, though, what would happen if we weaken this still further: what can we say if we drop the assumption that $\mathcal{K}$ is closed under directed colimits, and merely assume that the colimits that exist in $\mathcal{K}$ and are “**Set**-like” are those that are μ -directed for some μ ?

Here we introduce a framework, called *μ -abstract elementary classes*, that represents a model-theoretic approximation of that generalized notion, and which, most importantly, encompasses all of the examples considered in this introduction, including classes of models in infinitary logics $L_{\kappa,\mu}$, AECs, and μ -concrete AECs. This is not done just for the sake of generalization but in order to be able to deal with specific classes of structures that allow functions with infinite arity (like σ -complete Boolean algebras or formal power series). Moreover, such classes also occur naturally in the development of the classification theory for AECs, as can be seen by their use in [Vas] and [BoVaa] (there the class studied is the μ -AEC of μ -saturated models of an AEC).

We define μ -AECs in Section 2. We then show that the examples discussed above fit into this framework. We establish an analog of Shelah’s Presentation Theorem for μ -AECs in Section 3. In Section 4 we show that μ -AECs are, in fact, extraordinarily general: up to equivalence of categories, the μ -AECs are precisely the accessible categories whose morphisms are monomorphisms. Although this presents certain obstacles—it follows immediately that a general μ -AEC will not admit Ehrenfeucht-Mostowski constructions—there is a great deal that can be done on the μ -AEC side of this equivalence. In section 5, we show assuming the existence of large cardinals that μ -AECs satisfy tameness, an important locality property in the study of AECs. In Section 6, with the additional assumption of directed bounds, we begin to develop the classification theory of μ -AECs. Note that the results of Sections 5 and 6 transfer immediately to accessible categories with monomorphisms.

This paper was written while the fifth author was working on a Ph.D. thesis under the direction of the second author at Carnegie Mellon University and he would like to thank Professor Grossberg for his guidance and assistance in his research in general and in this work specifically. We also thank the referee for questions that helped us clarify some aspects of this paper.

2. PRELIMINARIES

We now introduce the notion of a μ -abstract elementary class, or μ -AEC. As with ordinary AECs, we give a semantic/axiomatic definition for a class of structures and a notion of strong substructure.

Definition 2.1. Fix an infinite cardinal μ .

A μ -ary language L consists of a set of function symbols $\langle F_i : i \in I_F \rangle$ and relations $\langle R_j : j \in J_R \rangle$ (here, I_F, J_R are index sets) so that each symbol has an arity, denoted $n(F_i)$ or $n(R_j)$, where n is an ordinal valued function $n : \{R_j, F_i \mid i \in I_F, j \in J_R\} \rightarrow \mu$.

Given a μ -ary language L , an L -structure M is $\langle |M|, F_i^M, R_j^M \rangle_{i \in I_F, j \in J_R}$ where $|M|$ is a set, called the universe of M ; $F_i^M : {}^{n(F_i)}|M| \rightarrow |M|$ is a function of arity $n(F_i)$; and $R_j^M \subset {}^{n(R_j)}|M|$ is a relation of arity $n(R_j)$.

We say that $\langle \mathcal{K}, \leq_{\mathcal{K}} \rangle$ is a μ -abstract class provided

- (1) $\mathcal{K}$ is a class of L -structure for a fixed μ -ary language $L := L(\mathcal{K})$.
- (2) $\langle \mathcal{K}, \leq_{\mathcal{K}} \rangle$ is a partially pre-ordered class (that is, $\leq_{\mathcal{K}}$ is reflexive and transitive) such that $M \leq_{\mathcal{K}} N$ implies that M is an L -submodel of N .
- (3) $\langle \mathcal{K}, \leq_{\mathcal{K}} \rangle$ respects L -isomorphisms; that is, if $f : N \rightarrow N'$ is an L -isomorphism and $N \in \mathcal{K}$, then $N' \in \mathcal{K}$ and if we also have $M \in \mathcal{K}$ with $M \leq_{\mathcal{K}} N$, then $f(M) \in \mathcal{K}$ and $f(M) \leq_{\mathcal{K}} N'$;

We often do not make the distinction between $\mathcal{K}$ and $(\mathcal{K}, \leq_{\mathcal{K}})$.

An L -homomorphism is called a substructure embedding if it is injective and reflects all relations. Both inclusions of a substructure and isomorphisms are substructure embeddings. Conversely, if $h : M \rightarrow N$ is a substructure embedding then M is isomorphic to the substructure $h(M)$ of N . The category of all L -structures and substructure embeddings is denoted by $\mathbf{Emb}(L)$. Then an abstract class is the same as a subcategory $\mathcal{K}$ of $\mathbf{Emb}(L)$ which is

- (1) *Replete*, i.e., closed under isomorphic objects.
- (2) *Iso-full*, i.e., containing isomorphisms between $\mathcal{K}$ -objects.

Let $\langle I, \leq \rangle$ be a partially ordered set. We say that I is μ -directed, where μ is a regular cardinal, provided that for every $J \subseteq I$ with $\text{card}(J) < \mu$, there exists $r \in I$ such that $r \geq s$ for all $s \in J$. Thus $\aleph_0$ -directed is the usual notion of directed set. Let $\langle \mathcal{K}, \leq_{\mathcal{K}} \rangle$ be an abstract class. A family $\{M_i \mid i \in I\} \subseteq \mathcal{K}$ is called a μ -directed system provided I is a μ -directed set and $i < j$ implies $M_i \leq_{\mathcal{K}} M_j$. This is the same as a μ -directed diagram in $\mathcal{K}$.

Definition 2.2. Suppose $\langle \mathcal{K}, \leq_{\mathcal{K}} \rangle$ is a μ -abstract class, with μ a regular cardinal. We say that $\langle \mathcal{K}, \leq_{\mathcal{K}} \rangle$ is a μ -abstract elementary class if the following properties hold:

- (1) (Coherence) *if $M_0, M_1, M_2 \in \mathcal{K}$ with $M_0 \leq_{\mathcal{K}} M_2$, $M_1 \leq_{\mathcal{K}} M_2$, and $M_0 \subseteq M_1$, then $M_0 \leq_{\mathcal{K}} M_1$;*
- (2) (Tarski-Vaught chain axioms) *If $\{M_i \in \mathcal{K} : i \in I\}$ is a μ -directed system, then:*
 - (a) $\bigcup_{i \in I} M_i \in \mathcal{K}$ *and, for all $j \in I$, we have $M_j \leq_{\mathcal{K}} \bigcup_{i \in I} M_i$; and*
 - (b) *if there is some $N \in \mathcal{K}$ so that, for all $i \in I$, we have $M_i \leq_{\mathcal{K}} N$, then we also have $\bigcup_{i \in I} M_i \leq_{\mathcal{K}} N$.*
- (3) (Löwenheim-Skolem-Tarski number axiom) *There exists a cardinal $\lambda = \lambda^{<\mu} \geq \text{card}(\mathbf{L}(\mathcal{K})) + \mu$ such that for any $M \in \mathcal{K}$ and $A \subseteq |M|$, there is some $N \leq_{\mathcal{K}} M$ such that $A \subseteq |N|$ and $\text{card}(N) \leq \text{card}(A)^{<\mu} + \lambda$. $\text{LS}(\mathcal{K})$ is the minimal cardinal λ with this property.*¹

Note that this definition mimics the definition of an AEC. We highlight the key differences:

Remark 2.3.

- (1) *Functions and relations are permitted to have infinite arity.*
- (2) *The Löwenheim-Skolem-Tarski axiom only guarantees the existence of submodels of certain cardinalities, subject to favorable cardinal arithmetic.*
- (3) *Closure under unions of $\leq_{\mathcal{K}}$ -increasing chains does not hold unconditionally: the directed systems must in fact be μ -directed.*
- (4) *The Tarski-Vaught axioms describe μ -directed systems rather than chains and say that $\mathcal{K}$ is closed under μ -directed colimits in $\mathbf{Emb}(\mathbf{L})$. One could have only required that every chain of models indexed by an ordinal of cofinality at least μ has a least upper bound. When $\mu = \aleph_0$, this is well-known to give an equivalent definition (see e.g. [AdRo94] 1.7, though the central idea of the proof dates back to Iwamura's Lemma, [Iwa44]). In general, though, this is significantly weaker (see [AdRo94] Exercise 1.c). Concretely, proving the presentation theorem becomes problematic if one opts instead for the chain definition.*
- (5) *Replacing the Tarski-Vaught axioms by $\mathcal{K}$ being closed under μ -directed colimits in $\mathbf{Emb}(\mathbf{L})$ makes sense also when $\mathcal{K}$ is a λ -abstract class for $\lambda > \mu$.*
- (6) *As a notational remark, we use $\text{card}()$ to denote the size of sets and (universes) of models. This breaks with convention, but is to avoid $|\cdot|$ being used to denote both universe and cardinality, leading to the notation $\|M\|$ for the cardinality of the universe of a model.*

As promised in the introduction, μ -AECs subsume many previously studied model theoretic frameworks:

¹Note that $\text{LS}(\mathcal{K})$ really depends on μ but μ will always be clear from context.

- (1) All AECs are $\aleph_0$ -AECs with the same Löwenheim-Skolem number. This follows directly from the definition. See [Gro02] for examples of classes of structures that are AECs; in particular, AECs subsume classical first-order model theory.
- (2) Given an AEC $\mathcal{K}$ with amalgamation (such as models of a first order theory), and a regular cardinal $\mu > \text{LS}(\mathcal{K})$, the class of μ -saturated² models of $\mathcal{K}$ is a μ -AEC with Löwenheim-Skolem number $\text{LS}(\mathcal{K})^{<\mu}$. If $\mathcal{K}$ is also tame and stable (or superstable), then the results of Boney and Vasey [Bo-Vaa] show that this is true even for certain cardinals below the saturation cardinal μ .
- (3) Let $\lambda \geq \mu$ be cardinals with μ regular. Let L_A be a fragment of $L_{\lambda,\mu}$ (recall that a fragment is a collection of formulas closed under sub formulas and first order connectives), and let T be a theory in that fragment. Then $\mathcal{K} = (\text{Mod } T, \preceq_{L_A})$ is a μ -AEC with $\text{LS}(\mathcal{K}) = (\text{card}(L_A) + \text{card}(T))^{<\mu}$, where $M \preceq_{L_A} N$ if and only if for all $\phi(\mathbf{x}) \in L_A$ and $\mathbf{m} \in |M|$ of matching arity (which might be infinite), we have that $M \models \phi[\mathbf{a}]$ if and only if $N \models \phi[\mathbf{a}]$.
- (4) Complete metric spaces form an $\aleph_1$ -AECs. This follows from the above item because metric spaces are axiomatizable in first order and completeness is axiomatized by the L_{ω_1,ω_1} sentence

$$\forall \langle x_n : n < \omega \rangle [(\bigwedge_{\epsilon \in \mathbb{Q}^+} \bigvee_{N < \omega} \bigwedge_{N < n < m < \omega} d(x_n, x_m) < \epsilon) \implies \exists y (\bigwedge_{\epsilon \in \mathbb{Q}^+} \bigvee_{N < \omega} \bigwedge_{N < n < \omega} d(x_n, y) < \epsilon)]$$

Although this does not capture the $[0, 1]$ -value nature of many treatments of the model theory of metric structures, such as [BBHU08], this can be incorporated in one of two ways. One could add the real numbers as a second sort, interpret relations as functions between the sorts, and axiomatize all of the continuity properties. A less direct approach is taken in [Bon], where a complete structure is approximated by a dense subset describable in $L_{\omega_1,\omega}$.

- (5) Along the lines of complete metric spaces, μ -complete boolean algebras are μ -AECs because μ -completeness can be written as a $L_{\mu,\mu}$ -sentence.
- (6) Any μ -concrete AEC (or μ -CAEC), in the sense of [LiRoa], is a μ -AEC.
- (7) Any μ -ary functorial expansion of a μ -AEC is naturally a μ -AEC. See Section 2.1 immediately below.
- (8) Generalizing $L(Q)$, consider classes axiomatized by $L_{\lambda,\mu}(Q^\chi)$, where Q^χ is the quantifier “there exist at least χ ” (the standard $L(Q)$ is $L_{\omega,\omega}(Q^{\aleph_1})$ in this notation). As in (3), let T be a theory in $L_{\lambda,\mu}(Q^\chi)$ and L_A be a fragment of this logic containing T . Since Q^χ is $L_{\chi,\chi}$ expressible, we already have $\mathcal{K}_0 := (\text{Mod } T, \preceq_{L_A})$ is a $(\mu + \chi)$ -AEC with $\text{LS}(\mathcal{K}_0) = (\text{card}(L_A) + \chi)^{<(\mu+\chi)}$. For a stronger result, if we set $M \preceq_{L_A}^* N$ by

²In the sense of Galois types.

$M \leq_{L_A} N$ and if $Q^x x \phi(x, \mathbf{y}) \in L_A$ with $M \models \neg Q^x x \phi(x, \mathbf{m})$, then
 $\phi(M, \mathbf{a}) = \phi(N, \mathbf{a})$

then $\mathcal{K}_1 := (\text{Mod } T, \leq_{L_A}^*)$ is a μ -AEC with $\text{LS}(\mathcal{K}_1) = (\text{card}(L_A) + \chi)^{<\mu}$.
 Moreover, if $\chi = \chi_0^+$ and L_A only contains *negative* instances of Q^x , then
 $\text{LS}(\mathcal{K}_1) = (\text{card}(L_A) + \chi_0)^{<\mu}$.

We now briefly discuss the interplay between certain μ -AECs and functorial expansions.

2.1. Functorial expansions and infinite summation. Recall from [Vasb, Definition 3.1]:

Definition 2.4. Let $\mathcal{K}$ be a μ -AEC with $L = L(\mathcal{K})$ and let $\widehat{L}$ be a λ -ary expansion of L with $\lambda \geq \mu$. A λ -ary $\widehat{L}$ -functorial expansion of $\mathcal{K}$ is a class $\widehat{\mathcal{K}}$ of $\widehat{L}$ -structures satisfying:

- (1) The map $\widehat{M} \mapsto \widehat{M} \upharpoonright L$ is a bijection from $\widehat{\mathcal{K}}$ onto $\mathcal{K}$. For $M \in \mathcal{K}$, we write $\widehat{M}$ for the unique element of $\widehat{\mathcal{K}}$ whose reduct is M .
- (2) Invariance: If $f : M \cong N$, then $f : \widehat{M} \cong \widehat{N}$.
- (3) Monotonicity: If $M \leq_{\mathcal{K}} N$, then $\widehat{M} \subseteq \widehat{N}$.

We order $\widehat{\mathcal{K}}$ by $\widehat{M} \leq_{\widehat{\mathcal{K}}} \widehat{N}$ if and only if $M \leq_{\mathcal{K}} N$.

Fact 2.5 (Proposition 3.8.(4) in [Vasb]). Let $\mathcal{K}$ be a μ -AEC and let $\widehat{\mathcal{K}}$ be a μ -ary functorial expansion of $\mathcal{K}$. Then $(\widehat{\mathcal{K}}, \leq_{\widehat{\mathcal{K}}})$ is a μ -AEC with $\text{LS}(\widehat{\mathcal{K}}) = \text{LS}(\mathcal{K})$.

Remark 2.6. A word of warning: if $\mathcal{K}$ is an AEC and $\widehat{\mathcal{K}}$ is a functorial expansion of $\mathcal{K}$, then $\mathcal{K}$ and $\widehat{\mathcal{K}}$ are isomorphic (as categories). In particular, any directed system in $\widehat{\mathcal{K}}$ has a colimit. However, $\widehat{\mathcal{K}}$ may not be an AEC if $L(\widehat{\mathcal{K}})$ is not finitary: the colimit of a directed system in $\widehat{\mathcal{K}}$ may not be the union: relations may need to contain more elements. However, if we change the definition of AEC to allow languages of infinite arity (see Remark 2.3.(5)), then $\widehat{\mathcal{K}}$ will be an AEC in that new sense, i.e. an “infinitary” AEC.

Remark 2.7. Let $\mathcal{K}$ be a μ -AEC and consider a $\widehat{L}$ -functorial expansion $\widehat{\mathcal{K}}$ of $\mathcal{K}$. Then any function and relation symbols from $\widehat{L}$ are interpretable in $\mathcal{K}$ in the sense of [Ros81] (this idea goes back to [Law63]). This means that function symbols of arity α are natural transformations $\varphi : U^\alpha \rightarrow U$ where $U : \mathcal{K} \rightarrow \mathbf{Set}$ is the forgetful functor (given as the domain restriction of the forgetful functor $\mathbf{Emb}(L) \rightarrow \mathbf{Set}$ assigning underlying sets to L -structures) and U^α is the functor $\mathbf{Set}(\alpha, U(-)) : \mathcal{K} \rightarrow \mathbf{Set}$. Similarly, relation symbols of arity α are subfunctors R of U^α .

If $\widehat{L}$ is μ -ary then subfunctors R preserve μ -directed colimits. Since $\mathcal{K}$ is an $\text{LS}(\mathcal{K})^+$ -accessible category (see 4.3), both φ and R are determined by their restrictions to the full subcategory $\mathcal{K}_{\text{LS}(\mathcal{K})^+}$ of $\mathcal{K}$ consisting of $\text{LS}(\mathcal{K})^+$ -presentable objects. Since there is only a set of such objects, there is a largest μ -ary functorial expansion where $\widehat{L}$ consists of all symbols for natural transformations and subfunctors as above. For $\mu = \aleph_0$, this is contained in [LiRo] Remark 3.5.

The main example in [Vasb] is Galois Morleyization (Definition 3.3 there). However there are many other examples including the original motivation for defining μ -AECs: infinite sums in boolean algebras. The point is that even though the language of boolean algebras with a sum operator is infinitary, we really need only to work in an appropriate class in a finitary language that we functorially expand as needed. This shows in a precise sense that the infinite sum operator is already implicit in the (finitary) structure of boolean algebras themselves.

Definition 2.8. Fix infinite cardinals $\lambda \geq \mu$. Let Φ be a set of formulas in $L_{\lambda, \mu}$. Let $(\mathcal{K}, \leq_{\mathcal{K}})$ be an abstract class. Define $\mathcal{K}_{\Phi} := (\mathcal{K}, \leq_{\mathcal{K}_{\Phi}})$ by $M \leq_{\mathcal{K}_{\Phi}} N$ if and only if $M \leq_{\mathcal{K}} N$ and $M \preceq_{\Phi} N$.

Lemma 2.9. Let $\lambda \geq \mu$, Φ be a set of formulas in $L_{\lambda, \mu}$. Let $(\mathcal{K}, \leq_{\mathcal{K}})$ be a μ_0 -AEC with $\mu_0 \leq \mu$. Then:

- (1) $\mathcal{K}_{\Phi}$ is a μ -AEC.
- (2) If all the formulas in Φ have fewer than μ_0 -many quantifiers, then $\mathcal{K}_{\Phi}$ satisfies the first Tarski-Vaught chain axiom of μ_0 -AECs.

Proof. The first part is straightforward. The second is proven by induction on the quantifier-depth of the formulas in Φ . $\square$

Example 2.10. Let T be a completion of the first-order theory of boolean algebras and let $\mathcal{K} := (\text{Mod}(T), \preceq)$. Let Φ consist of the L_{ω_1, ω_1} formula $\phi(\mathbf{x}, y)$ saying that y is a least upper bound of $\mathbf{x}$ (here $\ell(\mathbf{x}) = \omega$). Then ϕ has only one universal quantifier so by the Lemma, $\mathcal{K}_{\Phi}$ satisfies the first Tarski-Vaught chain axiom of AECs. Of course, $\mathcal{K}_{\Phi}$ is also an $\aleph_1$ -AEC. Now expand each $M \in \mathcal{K}$ to $\widehat{M}$ by defining $R_{\Sigma}^M(\mathbf{a}, b)$ to hold if and only if b is a least upper bound of $\mathbf{a}$ (with $\ell(\mathbf{a}) = \omega$). Let $\widehat{\mathcal{K}}_{\Phi} := \{\widehat{M} \mid M \in \mathcal{K}\}$. Then one can check that $\widehat{\mathcal{K}}_{\Phi}$ is a functorial expansion of $\mathcal{K}_{\Phi}$.

Basic definitions and concepts for AECs, such as amalgamation or Galois types (see [Bal09] or [Gro1X] for details), can be easily transferred to μ -AECs. In the following sections, we begin the process of translating essential theorems from AECs to μ -AECs.

3. PRESENTATION THEOREM

We now turn to the Presentation Theorem for μ -AECs. This theorem has its conceptual roots in Chang's Presentation Theorem [Cha68], which shows that $L_{\lambda, \omega}$ can be captured in a larger finitary language by omitting a set of types. A more immediate predecessor is Shelah's Presentation Theorem, which reaches the same conclusion for an arbitrary AEC. Unfortunately, while Chang's Presentation Theorem gives some insight into the original class, Shelah's theorem does not. However, the presentation is still a useful tool for some arguments and provides a syntactic characterization of what are otherwise purely semantic objects.

Definition 3.1. Let $L \subset L_1$ be μ -ary languages, T_1 an $(L_1)_{\mu, \mu}$ theory, and Γ be a set of μ -ary $(L_1)_{\mu, \mu}$ -types. Here we define a μ -ary $(L_1)_{\mu, \mu}$ -type as a set of $(L_1)_{\mu, \mu}$ formulas in the same free variables $\mathbf{x}$, where $\mathbf{x}$ has arity less than μ . We define

- $\text{EC}^\mu(T_1, \Gamma) = \{M : M \text{ an } L_1\text{-structure, } M \models T_1, M \text{ omits each type in } \Gamma\}$
- $\text{PC}^\mu(T_1, \Gamma, L) = \{M \upharpoonright L : M \in \text{EC}^\mu(T_1, \Gamma)\}$

Theorem 3.2. Let $\mathcal{K}$ be a μ -AEC with $\text{LS}(\mathcal{K}) = \chi$. Then we can find some $L_1 \supset L(\mathcal{K})$, a $(L_1)_{\mu, \mu}$ -theory T_1 of size χ , and a set Γ of μ -ary $(L_1)_{\mu, \mu}$ -types with $\text{card}(\Gamma) \leq 2^\chi$ so that $\mathcal{K} = \text{PC}^\mu(T_1, \Gamma, L(\mathcal{K}))$.

Although we don't state them here, the traditional moreover clauses (see e.g. the statement of [Bal09, Theorem 4.15]) apply as well.

Proof. We adapt the standard proofs; see, for instance, [Bal09, Theorem 4.15]. Set $\chi := \text{LS}(\mathcal{K})$. We introduce "Skolem functions" $L_1 := L(\mathcal{K}) \cup \{F_i^\alpha(\mathbf{x}) : i < \chi, \ell(\mathbf{x}) = \alpha < \mu\}$ and make very minimal demands by setting

$$T_1 := \{\exists x(x = x)\} \cup \{\forall \mathbf{x} F_i^\alpha(\mathbf{x}) = x_i : i < \alpha < \mu\}$$

For any $M_1 \models T_1$ and $\mathbf{a} \in |M_1|$, we define $N_{\mathbf{a}}^{M_1}$ to be the minimal L_1 -substructure of M_1 that contains $\mathbf{a}$. We can code the information about $N_{\mathbf{a}}^{M_1}$ into $\mathbf{a}$'s quantifier-free type:

$$p_{\mathbf{a}}^{M_1} := \{\phi(\mathbf{x}) : \phi(\mathbf{x}) \in (L_1)_{\mu, \mu} \text{ is quantifier-free and } M_1 \models \phi(\mathbf{a})\}$$

Given tuples $\mathbf{a} \in M_1$ and $\mathbf{b} \in M_2$ of the same length, we have that $p_{\mathbf{a}}^{M_1} = p_{\mathbf{b}}^{M_2}$ if and only if the map taking $\mathbf{a}$ to $\mathbf{b}$ induces an isomorphism $N_{\mathbf{a}}^{M_1} \cong N_{\mathbf{b}}^{M_2}$. Since we have this tight connection between types and structures, we precisely want to exclude types that give rise to structures not coming from $\mathcal{K}$. Thus, we set

$$\begin{aligned} \Gamma^{M_1} &:= \bigcup \{p_{\mathbf{a}}^{M_1} : \exists \mathbf{b} \subset \mathbf{a} \text{ such that } (N_{\mathbf{b}}^{M_1}) \upharpoonright L(\mathcal{K}) \not\leq_{\mathcal{K}} (N_{\mathbf{a}}^{M_1}) \upharpoonright L(\mathcal{K})\} \\ \Gamma &:= \bigcup_{M_1 \models T_1} \Gamma^{M_1} \end{aligned}$$

Note that the $\mathbf{b}$ in the first line might be $\mathbf{a}$, in which case the condition becomes $N_{\mathbf{a}}^{M_1} \upharpoonright L(\mathcal{K}) \notin \mathcal{K}$. By counting the number of $(L_1)_{\mu,\mu}$ -types, we have that $\text{card}(\Gamma) \leq 2^\chi$. Now all we have left to show is the following claim.

Claim: $\mathcal{K} = \text{PC}(T_1, \Gamma, L(\mathcal{K}))$

First, let $M_1 \in \text{EC}(T_1, \Gamma)$. Given $\mathbf{a} \in {}^{<\mu}M_1$, we know that $\mathbf{a} \models p_{\mathbf{a}}^{M_1}$ so $p_{\mathbf{a}}^{M_1} \notin \Gamma$. Thus, $\{N_{\mathbf{a}}^{M_1} \upharpoonright L(\mathcal{K}) : \mathbf{a} \in {}^{<\mu}M_1\}$ is a μ -directed system from $\mathcal{K}_{\leq \chi}$ with union $M_1 \upharpoonright L(\mathcal{K})$, so $M_1 \upharpoonright L(\mathcal{K}) \in \mathcal{K}$.

Second, let $M \in \mathcal{K}$. We need to define an expansion $M_1 \in \text{EC}(T_1, \Gamma)$. We can build a directed system $\{M_{\mathbf{a}} \in \mathcal{K}_{\chi} : \mathbf{a} \in {}^{<\mu}M\}$. Since each $M_{\mathbf{a}}$ has size χ , we can define the F_i^{α} by enumerating $|M_{\mathbf{a}}| = \{F_i^{\ell(\mathbf{a})}(\mathbf{a}) : i < \chi\}$ with the condition that $F_i^{\ell(\mathbf{a})}(\mathbf{a}) = a_i$ for $i < \ell(\mathbf{a})$. This precisely defines the expansion $M_1 := \langle M, F_i^{\alpha} \rangle_{i < \chi, \alpha < \mu}$. It is easy to see $M_1 \models T_1$. We also have $N_{\mathbf{a}}^{M_1} \upharpoonright L(\mathcal{K}) = M_{\mathbf{a}}$, so M_1 omits Γ because $\{M_{\mathbf{a}} : \mathbf{a} \in {}^{<\mu}M\}$ is a μ -directed system from $\mathcal{K}$. So $M \in \text{PC}^{\mu}(T_1, \Gamma, L)$. $\square$

Remark 3.3. A consequence of the presentation theorem for AECs is that an AEC $\mathcal{K}$ with a model of size $\beth_{(2^{\text{LS}(\mathcal{K})})^+}$ has arbitrarily large models (see e.g. [Bal09, Corollary 4.26]). The lack of Hanf numbers for $L_{\mu,\mu}$ means that we cannot use this to get similar results for μ -AECs. Thus the following question is still open: Can we compute a bound for the Hanf number $H(\lambda, \mu)$, where any μ -AEC $\mathcal{K}$ with $\text{LS}(\mathcal{K}) \leq \lambda$ that has a model larger than $H(\lambda, \mu)$ has arbitrarily large models?

4. μ -AECS AND ACCESSIBLE CATEGORIES

Accessible categories were introduced in [MaPa89] as categories closely connected with categories of models of $L_{\kappa,\lambda}$ theories. Roughly speaking, an accessible category is one that is closed under certain directed colimits, and whose objects can be built via certain directed colimits of a set of small objects. To be precise, we say that a category $\mathcal{K}$ is λ -accessible, λ a regular cardinal, if it is closed under λ -directed colimits (i.e. colimits indexed by a λ -directed poset) and contains, up to isomorphism, a set $\mathcal{A}$ of λ -presentable objects such that each object of $\mathcal{K}$ is a λ -directed colimit of objects from $\mathcal{A}$. Here λ -presentability functions as a notion of size that makes sense in a general, i.e. non-concrete, category: we say an object M is λ -presentable if its hom-functor $\mathcal{K}(M, -) : \mathcal{K} \rightarrow \mathbf{Set}$ preserves λ -directed colimits. Put another way, M is λ -presentable if for any morphism $f : M \rightarrow N$ with N a λ -directed colimit $\langle \phi_{\alpha} : N_{\alpha} \rightarrow N \rangle$, f factors essentially uniquely through one of the N_{α} , i.e. $f = \phi_{\alpha} f_{\alpha}$ for some $f_{\alpha} : M \rightarrow N_{\alpha}$.

For each regular cardinal κ , an accessible category $\mathcal{K}$ contains, up to isomorphism, only a set of κ -presentable objects. Any object M of a λ -accessible category is κ -presentable for some regular cardinal κ . Given an object M , the smallest cardinal κ such that M is κ -presentable is called the *presentability rank* of M . If the

presentability rank of M is a successor cardinal $\kappa = \|M\|^+$ then $\|M\|$ is called the *internal size* of M (this always happens if $\mathcal{K}$ has directed colimits or under GCH, see [BeRo12] 4.2 or 2.3.5). This notion of size internal to a particular category more closely resembles a notion of dimension—in the category **Met** of complete metric spaces with isometric embeddings, for example, the internal size of an object is precisely its density character—and, even in case the category is concrete, may not correspond to the cardinality of underlying sets. This distinction will resurface most clearly in the discussion at the beginning of Section 6 below.

We consider the category-theoretic structure of μ -AECs. As we will see, for any uncountable cardinal μ , any μ -AEC with Löwenheim-Skolem-Tarski number λ is a λ^+ -accessible category whose morphisms are monomorphisms, and that (perhaps more surprisingly) any μ -accessible category whose morphisms are monomorphisms is equivalent to a μ -AEC with Löwenheim-Skolem-Tarski number $\lambda = \max(\mu, \nu)^{<\mu}$, where ν , discussed in detail below, is the number of morphisms between μ -presentable objects.

It is of no small interest that a general μ -accessible category also satisfies a Löwenheim-Skolem-Tarski axiom of sorts, governed by the sharp inequality relation, $\trianglelefteq^3$. As we will see, this notion (see [MaPa89]) matches up perfectly with the behavior of μ -AECs conditioned by axiom 2.2(3).

We wish to show that μ -AECs and accessible categories are equivalent. For the easy direction—that every μ -AEC is accessible—we simply follow the argument for the corresponding fact for AECs in Section 4 of [Lie11].

Lemma 4.1. *Let $\mathcal{K}$ be a μ -AEC with Löwenheim-Skolem-Tarski number λ . Any $M \in \mathcal{K}$ can be expressed as a λ^+ -directed union of its $\leq_{\mathcal{K}}$ -substructures of size at most λ .*

Proof. Consider the diagram consisting of all $\leq_{\mathcal{K}}$ -substructures of M of size at most λ and with arrows the $\leq_{\mathcal{K}}$ -inclusions. To check that this diagram is λ^+ -directed, we must show that any collection of fewer than λ^+ many such submodels have a common extension also belonging to the diagram. Let $\{M_\alpha \mid \alpha < \nu\}$, $\nu < \lambda^+$, be such a collection. Since λ^+ is regular, $\sup\{|M_\alpha| \mid \alpha < \nu\} < \lambda^+$, whence

$$\text{card}\left(\bigcup_{\alpha < \nu} M_\alpha\right) \leq \nu \cdot \sup\{\text{card}(M_\alpha) \mid \alpha < \nu\} \leq \nu \cdot \lambda = \lambda$$

This set will be contained in some $M' \leq_{\mathcal{K}} M$ with $\text{card}(M') \leq \lambda^{<\nu} + \lambda = \lambda + \lambda = \lambda$, by the Löwenheim Skolem-Tarski axiom. For each $\alpha < \nu$, $M_\alpha \leq_{\mathcal{K}} M$ and $M_\alpha \subseteq M'$. Since $M' \leq_{\mathcal{K}} M$, coherence implies that $M_\alpha \leq_{\mathcal{K}} M'$. So we are done. $\square$

³The sharp inequality was introduced by Makkai and Pare [MaPa89, Section 2.3] and is defined by $\kappa \trianglelefteq \kappa'$ if and only if every κ -accessible category is also a κ' -accessible category, among other equivalent conditions.

Lemma 4.2. *Let $\mathcal{K}$ be a μ -AEC with Löwenheim-Skolem-Tarski number λ . A model M is λ^+ -presentable in $\mathcal{K}$ if and only if $\text{card}(M) \leq \lambda$.*

Proof. See the proof of Lemma 4.3 in [Lie11]. □

Taken together, these lemmas imply that any μ -AEC with Löwenheim-Skolem-Tarski number λ contains a set of λ^+ -presentable objects, namely $\mathcal{K}_{<\lambda^+}$, and that any model can be built as a λ^+ -directed colimit of such objects. As the Tarski-Vaught axioms ensure closure under μ -directed colimits and $\lambda \geq \mu$, it follows that $\mathcal{K}$ is closed under λ^+ -directed colimits. Thus we have:

Theorem 4.3. *Let $\mathcal{K}$ be a μ -AEC with Löwenheim-Skolem-Tarski number λ . Then $\mathcal{K}$ is a λ^+ -accessible category.*

Remark 4.4. *Theorem 4.3 is valid for any λ from 2.2(3) and not only for the minimal one. Moreover, we only need that λ satisfies the Löwenheim-Skolem-Tarski property for $\text{card}(A) \leq \lambda$. In this case, we will say that λ is a weak Löwenheim-Skolem-Tarski number.*

We now aim to prove that any accessible category whose morphisms are monomorphisms is a μ -AEC for some μ . In fact, there are two cases delineated below, concrete and abstract. In Theorem 4.5 we consider the concrete case: $\mathcal{K}$ is taken to be a κ -accessible category of L -structures and L -embeddings for some μ -ary language L where $\mu + \text{card}(L) \leq \kappa$. In particular, we insist that $\mathcal{K}$ sits nicely in $\mathbf{Emb}(L)$, the category of all L -structures and substructure embeddings—we may assume L is relational. In Theorem 4.10, we consider abstract accessible categories, with no prescribed signature or underlying sets.

Theorem 4.5. *Let L be a μ -ary signature and $\mathcal{K}$ be an iso-full, replete and coherent κ -accessible subcategory of $\mathbf{Emb}(L)$ where $\mu + \text{card}(L) \leq \kappa$. If $\mathcal{K}$ is closed under μ -directed colimits in $\mathbf{Emb}(L)$ and the embedding $\mathcal{K} \rightarrow \mathbf{Emb}(L)$ preserves κ -presentable objects then $\mathcal{K}$ is a μ -AEC with $\text{LS}(\mathcal{K}) \leq \lambda = \kappa^{<\mu}$.*

Proof. We verify that $\mathcal{K}$ satisfies the axioms in Definitions 2.1 and 2.2.

Given such a category, we define the relation $\leq_{\mathcal{K}}$ as we must: for $M, N \in \mathcal{K}$, $M \leq_{\mathcal{K}} N$ if and only if $M \subseteq N$ and the inclusion is a morphism in $\mathcal{K}$. Axiom 2.1(1) follows immediately from this definition. Axiom 2.1(2) follows from the assumption that the inclusion E is replete and iso-full, while 2.2(1) follows from the assumption that the aforementioned inclusion is a coherent functor. 2.2(2) is easily verified: given a μ -directed system $\langle M_i \mid i \in I \rangle$ in $\mathcal{K}$, the colimit lies in $\mathcal{K}$ (by μ -accessibility), and since the inclusion E preserves μ -directed colimits, it will be precisely the union of the system. So $\mathcal{K}$ is closed under μ -directed unions. The other clauses of 2.2(2) are clear as well.

Axiom 2.2(3), the Löwenheim-Skolem-Tarski Property, poses more of a challenge. To begin, we recall that in $\mathbf{Emb}(L)$, an object is κ^+ -presentable for $\kappa = \kappa^{<\mu} \geq \mu + \text{card}(L)$ precisely if its underlying set is of cardinality at most κ .

Recall that we intend to show that $\lambda = \kappa^{<\mu}$ satisfies 2.2(3). Let $M \in \mathcal{K}$ and $A \subseteq |M|$ with $|A| = \alpha > \lambda$. We begin by showing that $\mathcal{K}$ is $(\alpha^{<\mu})^+$ -accessible. This is a consequence of [LiRo94] 4.10 because $\kappa \leq \lambda < (\alpha^{<\mu})^+$ and $\mu \leq (\alpha^{<\mu})^+$. The sharp inequality is a consequence of Example 2.13(4) in [AdRo94]: for any cardinals $\beta < (\alpha^{<\mu})^+$ and $\gamma < \mu$,

$$\beta^\gamma \leq (\alpha^{<\mu})^\gamma = \alpha^{<\mu} < (\alpha^{<\mu})^+.$$

Since $\mathcal{K}$ is $(\alpha^{<\mu})^+$ -accessible, we can express M as an $(\alpha^{<\mu})^+$ -directed colimit of $(\alpha^{<\mu})^+$ -presentable objects in $\mathcal{K}$, say $\langle M_i \rightarrow M \mid i \in I \rangle$ —indeed, we may assume without loss that this is a $(\alpha^{<\mu})^+$ -directed system of inclusions. Following [LiRo94] 4.6, $E : \mathcal{K} \rightarrow \mathbf{Emb}(L)$ preserves $(\alpha^{<\mu})^+$ -presentable objects—hence the M_i are also $(\alpha^{<\mu})^+$ -presentable in $\mathbf{Emb}(L)$, and thus of cardinality at most $\alpha^{<\mu} = \text{card}(A)^{<\mu}$, by the remark in the previous paragraph. For each $a \in A$, choose M_{i_a} with $a \in |M_{i_a}|$. The set of all such M_{i_a} is of size at most $\alpha < (\alpha^{<\mu})^+$ and we have chosen the colimit to be $(\alpha^{<\mu})^+$ -directed, so there is some $M' = M_j$, $j \in I$, with $M_{i_a} \leq_{\mathcal{K}} M'$ for all $a \in A$. Hence $A \subseteq |M'|$, $M' \leq_{\mathcal{K}} M$, and

$$\text{card}(M') \leq \alpha^{<\mu} \leq \alpha^{<\mu} + \lambda = \text{card}(A)^{<\mu} + \lambda.$$

We now consider the case $\text{card}(A) \leq \lambda$. Hence

$$\text{card}(A)^{<\mu} \leq \lambda^{<\mu} = \lambda$$

So the cardinal bound in the Löwenheim-Skolem-Tarski Property defaults to λ . Since $\mu \leq \lambda^+$ (by [AdRo94] 2.13(4) again) and $\kappa \leq \lambda^+$, [LiRo94] 4.10 and 4.6 imply that $\mathcal{K}$ is λ^+ -accessible and the functor $E : \mathcal{K} \rightarrow \mathbf{Emb}(L)$ preserves λ^+ -presentable objects. Thus we may use the same argument as above to find $M' \leq_{\mathcal{K}} M$ of size λ containing A . $\square$

Remark 4.6. *Following 4.4 and 4.5, any μ -abstract class from 2.2 with (3) weakened to the existence of a weak Löwenheim-Skolem-Tarski number λ is a μ -AEC with $\text{LS}(\mathcal{K}) \leq (\lambda^+)^{<\mu}$.*

Assuming Vopěnka's principle, the weak Löwenheim-Skolem-Tarski number axiom is satisfied by any full subcategory $\mathcal{K}$ of $\mathbf{Emb}(L)$. This follows from [AdRo95] and is related to the unpublished theorem of Stavi (see [MV]).

To summarize, we have so far shown that any reasonably embedded κ -accessible subcategory of a category of structures $\mathbf{Emb}(L)$ is a μ -AEC. We wish to go further, however: given any μ -accessible category whose morphisms are monomorphisms, we claim that it is equivalent—as an abstract category—to a μ -AEC, in a sense that we now recall.

Definition 4.7. We say that categories $\mathcal{C}$ and $\mathcal{D}$ are equivalent if the following equivalent conditions (see [Mac71] V.4.1) hold:

- (1) There is a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ that is
 - full: For any C_1, C_2 in $\mathcal{C}$, the map $f \mapsto F(f)$ is a surjection from $\text{Hom}_{\mathcal{C}}(C_1, C_2)$ to $\text{Hom}_{\mathcal{D}}(FC_1, FC_2)$.
 - faithful: For any C_1, C_2 in $\mathcal{C}$, the map $f \mapsto F(f)$ is an injection from $\text{Hom}_{\mathcal{C}}(C_1, C_2)$ to $\text{Hom}_{\mathcal{D}}(FC_1, FC_2)$.
 - essentially surjective: Any object D in $\mathcal{D}$ is isomorphic to $F(C)$ for some C in $\mathcal{C}$.
- (2) There are functors $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$ such that the compositions FG and GF are naturally isomorphic to the identity functors on $\mathcal{D}$ and $\mathcal{C}$, respectively.

One might insist that the compositions in condition (2) are in fact equal to the identity functors, but this notion (*isomorphism of categories*) is typically too strong to be of interest—equivalence of categories as described above is sufficient to ensure that a pair of categories exhibit precisely the same properties. In particular, if $F : \mathcal{C} \rightarrow \mathcal{D}$ gives an equivalence of categories, it preserves and reflects internal sizes and gives a bijection between the isomorphism classes in $\mathcal{C}$ and those in $\mathcal{D}$; thus questions of, e.g., categoricity have identical answers in $\mathcal{C}$ and $\mathcal{D}$.

We proceed by constructing, for a general μ -accessible category $\mathcal{K}$ whose morphisms are monomorphisms, a full, faithful, essentially surjective functor from $\mathcal{K}$ to $\mathcal{K}'$, where $\mathcal{K}'$ is a μ -AEC. We begin by realizing a μ -accessible category as a category of structures.

Lemma 4.8. *Let $\mathcal{K}$ be a μ -accessible category whose morphisms are monomorphisms. There is a unary many-sorted signature L such that $\mathcal{K}$ is fully embedded to an equational variety in $\mathbf{Emb}(L)$.*

Moreover, this full embedding preserves μ -directed colimits.

Proof. Let $\mathcal{A}$ be the full subcategory of μ -presentable objects in $\mathcal{K}$ (technically, we want $\mathcal{A}$ to be skeletal, which makes it small). Consider the canonical embedding

$$E : \mathcal{K} \rightarrow \mathbf{Set}^{\mathcal{A}^{op}}$$

that takes each $K \in \mathcal{K}$ to the contravariant functor $\text{Hom}_{\mathcal{K}}(-, K)|_{\mathcal{A}^{op}}$, and each $\mathcal{K}$ -morphism $f : K \rightarrow K'$ to the natural transformation $E(f) : \text{Hom}_{\mathcal{K}}(-, K) \rightarrow \text{Hom}_{\mathcal{K}}(-, K')$ given by postcomposition with f . We note that, by Proposition 2.8 in [AdRo94], this functor is fully faithful and preserves μ -directed colimits. In fact, we may identify the image $\mathcal{L}_1$ of $\mathcal{K}$ in $\mathbf{Set}^{\mathcal{A}^{op}}$ with an equational variety. Let L be a signature with sorts $\{S_A \mid A \in \mathcal{A}\}$, and with unary function symbols for each morphism in $\mathcal{A}$, i.e. a function symbol $\bar{f}$ of arity $S_A \rightarrow S_B$ for each $\mathcal{A}$ -map $f : B \rightarrow A$, subject to certain equations: whenever $h = f \circ g$ in $\mathcal{A}$, we

insist that $\bar{h} = \bar{g}\bar{f}$. Concretely, the identification is given by a functor $F : \mathcal{L}_1 \rightarrow \mathbf{Emb}(L)$ that takes each functor $E(K) = \text{Hom}_{\mathcal{K}}(-, K)$ to the structure $FE(K)$ with sorts $S_A^{FE(K)} = \text{Hom}_{\mathcal{K}}(A, K)$ and with each $\bar{f} : S_A \rightarrow S_B$ interpreted as the function $\bar{f}^{FE(K)} : \text{Hom}_{\mathcal{K}}(A, K) \rightarrow \text{Hom}_{\mathcal{K}}(B, K)$ given by precomposition with f . Any morphism $g : K \rightarrow K'$ in $\mathcal{K}$ is first sent to the natural transformation $E(g) : \text{Hom}_{\mathcal{K}}(-, K) \rightarrow \text{Hom}_{\mathcal{K}}(-, K')$ then sent, via F , to $FE(g) : FE(K) \rightarrow FE(K')$, which is given sortwise by postcomposition with g , i.e. for any $A \in \mathcal{A}$ and $f \in S_A^{FE(K)} = \text{Hom}_{\mathcal{K}}(A, K)$, $g(f) = g \circ f$. Clearly, morphisms are injective in the image of $\mathcal{K}$ under FE , as they come from monomorphisms in $\mathcal{K}$, and they trivially reflect relations, meaning that in fact $F : \mathcal{L}_1 \rightarrow \mathbf{Emb}(L)$. $\square$

Let $\mathcal{L}_2$ denote the image of $\mathcal{K}$ in $\mathbf{Emb}(L)$ under FE . So we have exhibited $\mathcal{K}$ as a full subcategory of $\mathbf{Emb}(L)$ closed under μ -directed colimits, where Σ is a finitary language. As a result, the induced relation $\leq_{\mathcal{K}}$ is simply $\subseteq$, and iso-fullness and repleteness of the embedding are trivial. There is only one more wrinkle that we need to consider: the presentability rank of structures in the image of $\mathcal{K}$ in $\mathbf{Emb}(L)$ need not correspond to the cardinality of the union of their sorts—that is, if U denotes the forgetful functor $\mathbf{Emb}(L) \rightarrow \mathbf{Set}$, a μ -presentable object $FE(K)$ need not have $|UFE(K)| < \mu$ —so the argument in Theorem 4.5 cannot simply be repeated here. Still, U can only do so much damage:

Lemma 4.9. *The functor $U : \mathcal{L}_2 \rightarrow \mathbf{Set}$ sends μ -presentable objects to ν^+ -presentable objects, where $\nu = \text{card}(\text{Mor}(\mathcal{A}))$.*

Theorem 4.10. *Let $\mathcal{K}$ be a μ -accessible category with all morphisms mono. Then $\mathcal{K}$ is equivalent to a μ -AEC with Löwenheim-Skolem-Tarski number $\lambda = \max(\mu, \nu)^{<\mu}$.*

Proof. Consider $X \subseteq FE(K)$, and let $\alpha = \text{card}(X)$. Since $\mathcal{L}_2$ is μ -accessible, it is $(\alpha^{<\mu})^+$ -accessible (provided $\alpha \geq \mu$); see the proof of 4.5. Thus there is an $(\alpha^{<\mu})^+$ -presentable $\mathcal{L}_2$ -subobject M_X of $FE(K)$ with $X \subseteq M_X$. By Theorem 2.3.11 in [MaPa89], M_X can be expressed as an $(\alpha^{<\mu})^+$ -small μ -directed colimit of μ -presentables in $\mathcal{L}_2$, meaning that $U(M_X)$ is an $(\alpha^{<\mu})^+$ -small μ -directed colimit of sets of size less or equal than ν . This is of cardinality less or equal than $\alpha^{<\mu} + \max(\mu, \nu)$. This suggests $\max(\mu, \nu)$ might serve as our Löwenheim-Skolem-Tarski number, but we must fulfill the requirement that $\lambda^{<\mu} = \mu$. So, take $\lambda = \max(\mu, \nu)^{<\mu}$. $\square$

The μ -AEC from 4.10 is a full subcategory of $\mathbf{Emb}(L)$ where L is a finitary language. Although this equivalence destroys both the ambient language and the underlying sets, and thus moves beyond the methods usually entertained in model theory, it allows us to transfer intuition and concepts between the two contexts.

The equivalence allows us to generate the notion of a Löwenheim-Skolem number in a general accessible concrete category, where concreteness is necessary to form the question.

Proposition 4.11. *Let $(\mathcal{K}, U)$ be a μ -accessible concrete category with all maps monomorphisms such that U preserves μ -directed colimits. Then if $M \in \mathcal{K}$ and $X \subset UM$, there is a subobject $M_0 \in \mathcal{K}$ of M such that $X \subset UM_0$ and M_0 is $(\text{card}(X)^{<\mu})^+$ -presentable.*

Note that we have proved that every μ -accessible category with all maps monomorphism has such a concrete functor: by Theorem 4.10, there is a (full and faithful) equivalence $F : \mathcal{K} \rightarrow \mathcal{K}'$, where $\mathcal{K}'$ is some μ -AEC. The universe functor $U : \mathcal{K}' \rightarrow \mathbf{Set}$ is faithful and preserves μ -directed colimits, so $FU : \mathcal{K} \rightarrow \mathcal{K}'$ does as well.

Proof. Let $M \in \mathcal{K}$, where $\mathcal{K}$ is μ -accessible and concrete with monomorphisms. Let $X \subset UM$. We want to find $M_0 \leq M$ with $X \subset UM_0$ that is $(\text{card}(X)^{<\mu})^+$ -presentable. By accessibility, we can write M as a μ -directed colimit $\langle M^i, f_{j,i} \mid j < i \in I \rangle$ where $M^i \in \mathcal{K}$ is μ -presentable, I is μ -directed, and $f_{i,\infty}$ are the colimit maps.

Because U preserves μ -directed colimits, there is $I_0 \subset I$ of size $\leq \text{card}(X)$ such that, for every $x \in X$, there is some $i_x \in I_0$ such that $x \in Uf_{i_x,\infty}M^{i_x}$. Close this to a μ -directed subset $I_1 \subset I$ of size $\leq \text{card}(X)^{<\mu}$ and let $(M^*, f_{i,*})$ be the colimit of $\{M^i, f_{j,i} \mid j < i \in I_1\}$. Since this system also embeds into M , there is a canonical map $f^* : M^* \rightarrow M$. Set $M_0 = f^*M^*$. Then M_0 is a subobject of M and $X \subset UM_0$, so we just need to show M_0 is $(\text{card}(X)^{<\mu})^+$ -presentable. This follows from [AdRo94] 1.16: since $\mu \leq (\text{card}(X)^{<\mu})^+$, each M^i is $(\text{card}(X)^{<\mu})^+$ -presentable. Since $\text{card}(I_1) \leq \text{card}(X)^{<\mu}$, M_0 is $(\text{card}(X)^{<\mu})^+$ -presentable by the cited result. □

Although the previous theorem doesn't use any model theoretic properties directly, it is inspired by standard proofs of the downward Löwenheim-Skolem theorem and doesn't seem to be previously known.

Going the other direction, knowledge about accessible categories allows us to show that μ -AECs do not, in general, admit Ehrenfeucht-Mostowski constructions. In particular, not every μ -AEC $\mathcal{K}$ admits a faithful functor $E : \mathbf{Lin} \rightarrow \mathcal{K}$:

Example 4.12. Let $\mathcal{K}$ be the category of well-ordered sets and order-preserving injections. By [AdRo94] 2.3(8), $\mathcal{K}$ is ω_1 -accessible, and clearly all of its morphisms are monomorphisms. By Theorem 4.10, it is therefore equivalent to an ω_1 -AEC. As $\mathcal{K}$ is isomorphism rigid—that is, it contains no nonidentity isomorphisms—it cannot admit a faithful functor from $\mathbf{Lin}$, which is far from isomorphism rigid.

Ehrenfeucht-Mostowski constructions are a very powerful tool in the study of AECs (see for example [Sh394]). This suggests that μ -AECs may be too general to support a robust classification theory. In particular, the lack Ehrenfeucht-Mostowski models, in turn, means that there is no analogue of the Hanf number that has proven to be very useful in the study of AECs.

A possible substitute to the notion of Hanf number is that of *LS-accessibility*, which was introduced by [BeRo12]. Rather than looking at the cardinality of the models, they asked about the internal size, as computed in the category. The shift stems from the following: it is clear that there are $\aleph_1$ -AECs that don't have models in arbitrarily large cardinalities: looking at complete (non-discrete) metric spaces or [BeRo12, Example 4.8], there can be no models in cardinalities satisfying $\lambda < \lambda^\omega$. However, the internal size based on presentability rank mentioned above gives that, e. g., complete metric spaces have models of all *sizes*. Thus, an accessible category is called LS-accessible iff there is a threshold such that there are object of every *size* above that threshold. Beke and Rosicky [BeRo12] ask if every large accessible category is LS-accessible. This question is still open and a positive answer (even restricting to accessible categories where all maps are mono) would aid the analysis of μ -AECs (see the discussion at the start of Section 6).

Still, under the additional assumption of upper bounds for increasing chains of structures—*directed bounds*, in the language of [Ros97], or the δ -*chain extension property*, defined below—we can rule out Example 4.12, and begin to develop a genuine classification theory.

5. TAMENESS AND LARGE CARDINALS

In [Bon14], it was shown by the first author that, assuming the existence of large cardinals, every AEC satisfies the important locality property know as tameness. Tameness was isolated (from an argument of Shelah [Sh394]) by Grossberg and VanDieren in [GrVa06a], and was used to prove an upward categoricity transfer from a successor cardinal in [GrVa06b, GrVa06c]. Tame AECs have since been a very productive area of study. For example, they admit a well-behaved notion of independence [Vasc, Vasa] and many definitions of superstability can be shown to be equivalent in the tame context [GrVa].

In this section, we generalize Boney's theorem to μ -AECs (in a sense, this also partially generalizes the recent [BoZa] which proved an analogous result for metric AECs, but for a stronger, metric specialization of tameness). We start by recalling the definition of tameness (and its generalization: full tameness and shortness) to this context. This generalization already appears in [Vasb, Definition 2.21].

Definition 5.1 (Definitions 3.1 and 3.3 in [Bon14]). *Let $\mathcal{K}$ be an abstract class and let κ be an infinite cardinal.*

- (1) $\mathcal{K}$ is $(< \kappa)$ -tame if for any distinct $p, q \in gS(M)$, there exists $A \subseteq \text{card}(M)$ such that $|A| < \kappa$ and⁴ $p \restriction A \neq q \restriction A$.
- (2) $\mathcal{K}$ is fully $(< \kappa)$ -tame and short if for any distinct $p, q \in gS^\alpha(M)$, there exists $I \subseteq \alpha$ and $A \subseteq |M|$ such that $\text{card}(I) + \text{card}(A) < \kappa$ and $p^I \restriction A \neq q^I \restriction A$.
- (3) We say $\mathcal{K}$ is tame if it is $(< \kappa)$ -tame for some κ , similarly for fully tame and short.

Instead of strongly compact cardinals, we will (as in [BoUn] and [BTRo]) use *almost* strongly compact cardinals:

Definition 5.2. *An uncountable limit cardinal κ is almost strongly compact if for every $\mu < \kappa$, every κ -complete filter extends to a μ -complete ultrafilter.*

Note that the outline here follows the original model theoretic arguments of [Bon14]. The category theoretic arguments of [LiRo] and [BTRo] can also be used.

A minor variation of the proof of Łoś's theorem for $L_{\kappa, \kappa}$ (see [Dic75, Theorem 3.3.1]) gives:

Fact 5.3. *Let κ be an almost strongly compact cardinal. Let $\mu < \kappa$, let $(M_i)_{i \in I}$ be L -structures, and let U be a μ^+ -complete ultrafilter on I . Then for any formula $\phi \in L_{\mu, \mu}$, $\prod M_i \setminus U \models \phi[[f]_U]$ if and only if $M_i \models \phi[f(i)]$ for U -almost all $i \in I$.*

Using the presentation theorem, we obtain Łoś's theorem for μ -AECs:

Lemma 5.4. *Let $\mathcal{K}$ be a μ -AEC. Let $(M_i)_{i \in I}$ be models in $\mathcal{K}$ and let U be a $(2^{\text{LS}(\mathcal{K})})^+$ -complete ultrafilter on I . Then $\prod M_i \setminus U \in \mathcal{K}$.*

Proof sketch. Let $\mu := (2^{\text{LS}(\mathcal{K})})^+$. By the presentation theorem (Theorem 3.2), there exists a language $L' \supseteq L(\mathcal{K})$ and a sentence $\phi \in L'_{\mu, \mu}$ such that $\mathcal{K} = \text{Mod}(\phi) \restriction L = \mathcal{K}$. Now use Fact 5.3 together with the proof of [Bon14, Theorem 4.3]. $\square$

All the moreover clauses of [Bon14, Theorem 4.3] are also obtained, thus by the same proof as [Bon14, Theorem 4.5], we get:

Theorem 5.5. *Let $\mathcal{K}$ be a μ -AEC and let $\kappa > \text{LS}(\mathcal{K})$ be almost strongly compact. Then $\mathcal{K}$ is fully $(< \kappa)$ -tame and short.*

In particular, if there is a proper class of almost strongly compact cardinals, every μ -AEC is fully tame and short. Using the recent converse for the special case of AECs due to Boney and Unger [BoUn], we obtain also a converse in μ -AECs:

Theorem 5.6. *The following are equivalent:*

⁴We use here Galois types over sets as defined in [Vasb, Definition 2.16].

- (1) *For every μ , every μ -AEC is fully tame and short.*
- (2) *Every AEC is tame.*
- (3) *There exists a proper class of almost strongly compact cardinals.*

Proof. (1) implies (2) is because AECs are $\aleph_0$ -AECs. (2) implies (3) is [BoUn] and (3) implies (1) is Theorem 5.5. $\square$

6. ON CATEGORICAL μ -AECs

Here we show that some non-trivial theorems of classification theory for AECs transfer to μ -AECs and, by extension, accessible categories with monomorphisms. Most of the classification theory for AECs has been driven by Shelah's categoricity conjecture⁵. For an abstract class $\mathcal{K}$, we write $I(\lambda, \mathcal{K})$ for the number of pairwise non-isomorphic models of $\mathcal{K}$ of cardinality λ . An abstract class $\mathcal{K}$ is said to be *categorical in λ* if $I(\lambda, \mathcal{K}) = 1$. Inspired by Morley's categoricity theorem, Shelah conjectured:

Conjecture 6.1. *If an AEC is categorical in a high-enough cardinal, then it is categorical on a tail of cardinals.*

Naturally, one can ask the same question for both μ -AECs and accessible categories, where, following [Ros97], we say an accessible category is *categorical in λ* if it contains exactly one object of internal size λ (up to isomorphism). By shifting the question to these more general frameworks, of course, we make it more difficult to arrive at a positive answer. If the answer is negative, on the other hand, counterexamples should be more readily available in our contexts: if indeed the answer is negative, this would give us a bound on the level of generality at which the categoricity conjecture can hold.

Question 6.2. *If a large accessible category (whose morphisms are monomorphisms) is categorical in a high-enough cardinal, is it categorical on a tail of cardinals?*

A negative answer to the question of Beke and Rosicky from Section 4—an example of an large accessible category $\mathcal{K}$ with arbitrarily large gaps in internal sizes—would yield a negative answer to Question 6.2: as noted in [BeRo12] 6.3, it suffices to take the coproduct $\mathcal{K} \coprod \mathbf{Set}$. This adds exactly one isomorphism class to each size, resulting in a category that is (internally) categorical in arbitrarily high cardinals—the gaps of $\mathcal{K}$ —but also fails to be (internally) categorical in arbitrarily large cardinals. By taking injective mappings of sets, one can do the same for large accessible categories whose morphisms are monomorphisms. [BeRo12] and [LiRo] contain sufficient conditions for LS-accessibility: in particular, it is enough to add the assumption of the existence of arbitrary directed colimits (see [LiRo], 2.7).

⁵For more references and history, see the introduction of Shelah's book [Sh:h]

For μ -AECs, the natural formulation is in terms not of the internal size, but of the cardinality of underlying sets. Some adjustments have to be made, as a μ -AEC need not have a model of cardinality λ when $\lambda^{<\mu} > \lambda$, and thus eventual categoricity would fail more or less trivially.

Question 6.3. *If a μ -AEC is categorical in a high-enough cardinal λ with $\lambda = \lambda^{<\mu}$, is it categorical in all sufficiently high λ' such that $\lambda' = (\lambda')^{<\mu}$.*

For $\mu = \omega$, this question reduces to 6.1.

Remark 6.4. *We will show that a positive answer to Question 6.2, the internal version, implies, at the very least, a positive answer to Question 6.1. Let $\mathcal{K}$ be an AEC in a language L . Then $\mathcal{K}$ is an accessible category and, following [BeRo12] 4.1, 4.3 and 3.6, there is a regular cardinal κ such that $\mathcal{K}$ is κ -accessible and E preserves sizes $\lambda \geq \kappa$. We can assume that, in $\mathbf{Emb}(L)$, they coincide with cardinalities of underlying sets. Thus, any K_1, K_2 with sufficiently large and distinct $\text{card}(EK_1), \text{card}(EK_2)$ have distinct sizes $\text{card}(K_1), \text{card}(K_2)$ and thus K_1 and K_2 are not isomorphic.*

At present we do not know whether a positive answer to 6.2 implies a positive answer to 6.3.

Of course 6.2 is currently out of reach, as is 6.3. We are not sure about the truth value of either one: it is plausible that there are counterexamples. A possible starting point for 6.3 would be to use Theorem 5.5 to try to generalize [Bon14] to μ -AECs categorical in an appropriate successor above a strongly compact (see also [Sh1019], which proves some model-theoretic results for classes of models of $L_{\kappa,\kappa}$ with κ a strongly compact cardinal).

We show here that some facts which follow from categoricity in AECs also follow from categoricity in μ -AECs. As in [Ros97], which considers categoricity in accessible categories with directed bounds (and, ultimately, directed colimits), we have to add the following hypothesis:

Definition 6.5. *Let δ be an ordinal. An abstract class $\mathcal{K}$ has the δ -chain extension property if for every chain $\langle M_i : i < \delta \rangle$, there exists $M_\delta \in \mathcal{K}$ such that $M_i \leq_{\mathcal{K}} M_\delta$ for all $i < \delta$. We say that $\mathcal{K}$ has the chain extension property if it has the δ -chain extension property for every limit ordinal δ .*

Remark 6.6. *If $\mathcal{K}$ is a μ -AEC, then $\mathcal{K}$ has the chain extension property if and only if $\mathcal{K}$ has the δ -chain extension property for every limit $\delta < \mu$.*

Remark 6.7. *μ -CAECs have the chain extension property (recall the item (6) from the list of examples). Moreover, any μ -AEC naturally derived from⁶ an AEC (such as the class of μ -saturated models of an AEC) will have the chain extension property.*

⁶This can be made precise using the notion of a *skeleton*, see [Vasa, Definition 5.3].

We adapt Shelah's [Sh:h, Theorem IV.1.12.(1)] to μ -AECs:

Theorem 6.8. *Let $\mathcal{K}$ be a μ -AEC. Let $\lambda \geq \text{LS}(\mathcal{K})$ be such that $\lambda = \lambda^{<\mu}$ and $\mathcal{K}_\lambda$ has the δ -chain extension property for all limit $\delta < \lambda^+$. Assume $\mathcal{K}$ is categorical in λ . Let $M, N \in \mathcal{K}_{\geq \lambda}$. If $M \leq_{\mathcal{K}} N$, then $M \preceq_{L_{\infty, \mu}} N$.*

Notice that the cardinal arithmetic ($\lambda^{<\mu} = \lambda$) is a crucial simplifying assumption in the AEC version that Shelah later worked to remove (see [Sh:h, Section IV.2] and [BoVab]). It appears naturally here in the context of a μ -AEC, but note that the chain extension might guarantee the existence of models of intermediate sizes (i.e. in $\chi < \chi^{<\mu}$).

Proof of Theorem 6.8. We first assume that $M, N \in \mathcal{K}_\lambda$ and $M \leq_{\mathcal{K}} N$. Let $\phi(\mathbf{y})$ be an $L_{\infty, \mu}$ -formula with $\ell(\mathbf{y}) = \alpha < \mu$ and let $\mathbf{a} \in {}^\alpha M$. We show that $M \models \phi[\mathbf{a}]$ if and only if $N \models \phi[\mathbf{a}]$ by induction on the complexity of ϕ . If ϕ is atomic, this holds because $M \subseteq N$. If ϕ is a boolean combination of formulas of lower complexity, this is easy to check too. So assume that $\phi(\mathbf{y}) = \exists \mathbf{x} \psi(\mathbf{x}, \mathbf{y})$. If $M \models \phi[\mathbf{a}]$, then using induction we directly get that $N \models \phi[\mathbf{a}]$. Now assume $N \models \phi[\mathbf{a}]$, and let $\mathbf{b} \in {}^{<\mu} N$ be such that $N \models \psi[\mathbf{b}, \mathbf{a}]$.

We build an increasing chain $\langle M_i : i < \lambda^+ \rangle$ and $\langle f_i, g_i : i < \lambda^+ \rangle$ such that for all $i < \lambda^+$:

- (1) $M_i \in \mathcal{K}_\lambda$
- (2) If $\text{cf } i \geq \mu$, then $M_i = \bigcup_{j < i} M_j$.
- (3) $f_i : M \cong M_i$.
- (4) $g_i : N \cong M_{i+1}$.
- (5) $f_i \subseteq g_i$.

This is possible. If $i = 0$, let $M_0 := M$. For any i , given M_i , use categoricity to pick $f_i : M \cong M_i$ and extend it to $g_i : N \cong M_{i+1}$. If i is limit and $\text{cf } i \geq \mu$, take unions. If $\text{cf } i < \mu$, use the chain extension property to find $M_i \in \mathcal{K}_\lambda$ such that $M_j \leq_{\mathcal{K}} M'_i$ for all $j < i$.

This is enough. For each $i < \lambda^+$, let $\alpha(i)$ be the least $\alpha < \lambda^+$ such that $\text{ran}(f_i(\mathbf{a})) \subseteq |M_\alpha|$. Let $S := \{i < \lambda^+ \mid \text{cf } i \geq \mu\}$. Note that S is a stationary subset of λ^+ and the map $i \mapsto \alpha(i)$ is regressive on S . By Fodor's lemma, there exists a stationary $S_0 \subseteq S$ and $\alpha_0 < \lambda^+$ such that for any $i \in S_0$, $\alpha(i) = \alpha_0$, i.e. $\text{ran}(f_i(\mathbf{a})) \subseteq |M_{\alpha_0}|$. Now $\text{card}({}^{<\mu} M_{\alpha_0}) = \lambda^{<\mu} = \lambda$ and $|S_0| = \lambda^+$ so by the pigeonhole principle there exists $i < j$ in S_0 such that $f_i(\mathbf{a}) = f_j(\mathbf{a})$. Now, since $N \models \psi[\mathbf{b}, \mathbf{a}]$, we must have $M_{i+1} \models \psi[g_i(\mathbf{b}), g_i(\mathbf{a})]$. By the induction hypothesis, $M_j \models \psi[g_i(\mathbf{b}), g_i(\mathbf{a})]$. Thus $M_j \models \phi[g_i(\mathbf{a})]$. Since $f_i \subseteq g_i$, $g_i(\mathbf{a}) = f_i(\mathbf{a})$ so $M_j \models \phi[f_i(\mathbf{a})]$. Since $f_i(\mathbf{a}) = f_j(\mathbf{a})$, we have that $M_j \models \phi[f_j(\mathbf{a})]$. Applying f_j^{-1} to this equation, we obtain $M \models \phi[\mathbf{a}]$, as desired.

This proves the result in case $M, N \in \mathcal{K}_\lambda$. If $M, N \in \mathcal{K}_{\geq \lambda}$ and $M \leq_{\mathcal{K}} N$, then, as before, we can find a μ -directed system $\langle N_{\mathbf{a}} \in \mathcal{K}_\lambda : \mathbf{a} \in {}^{<\mu}N \rangle$ with colimit N such that $\mathbf{a} \in |N_{\mathbf{a}}|$ and, if $\mathbf{a} \in {}^{<\mu}M$, then $N_{\mathbf{a}} \leq_{\mathcal{K}} M$.

As before we prove by induction on $\phi \in L_{\infty, \mu}$ that $M \models \phi[\mathbf{a}]$ if and only if $N \models \phi[\mathbf{a}]$. The interesting case is when $\phi = \exists \mathbf{x} \psi(\mathbf{x}, \mathbf{y})$ and the left to right direction is straightforward, so assume $N \models \phi[\mathbf{a}]$, i. e., there exists $\mathbf{b} \in {}^{<\mu}|N|$ such that $N \models \psi[\mathbf{b}, \mathbf{a}]$. By the previous part, $N_{\mathbf{a}} \preceq_{L_{\infty, \mu}} N_{\mathbf{ab}}$. So there is $\mathbf{b}' \in N_{\mathbf{a}}$ such that $N_{\mathbf{a}} \models \psi[\mathbf{b}', \mathbf{a}]$. Since $N_{\mathbf{a}} \leq_{\mathcal{K}} M$, by induction, we have $M \models \phi[\mathbf{a}]$. $\square$

Another result that can be adapted is Shelah's famous combinatorial argument that amalgamation follows from categoricity in two successive cardinals [Sh88, Theorem 3.5]. We start with some simple definitions and lemmas:

Definition 6.9. Let $\mu \leq \lambda$ be regular cardinals. $C \subseteq \lambda$ is a μ -club if it is unbounded and whenever $\langle \alpha_i : i < \delta \rangle$ is increasing in C with $\mu \leq \text{cf } \delta < \lambda$, then $\sup_{i < \delta} \alpha_i \in C$.

Remark 6.10. So $\aleph_0$ -club is the usual notion of club.

Lemma 6.11. Let μ be a regular cardinal. Assume $\mathcal{K}$ is a μ -AEC and $\lambda \geq \text{LS}(\mathcal{K})$. Let $\langle M_i^\ell : i < \lambda^+ \rangle$, $\ell = 1, 2$, be increasing in $\mathcal{K}_\lambda$ such that for all $i < \lambda^+$ with $\text{cf } i \geq \mu$, $M_i^\ell = \bigcup_{j < i} M_j^\ell$.

If $f : \bigcup_{i < \lambda^+} M_i^1 \cong \bigcup_{i < \lambda^+} M_i^2$, then the set $\{i < \lambda^+ \mid f \upharpoonright M_i^1 : M_i^1 \cong M_i^2\}$ is a μ -club.

Proof. Let $C := \{i < \lambda^+ \mid f \upharpoonright M_i^1 : M_i^1 \cong M_i^2\}$. By cardinality considerations, for each $i < \lambda^+$, there is $j_i < \lambda^+$ such that $|f(M_i^1)| \subseteq |M_{j_i}^2|$ (by coherence this implies $f(M_i^1) \leq_{\mathcal{K}} M_{j_i}^2$). Let δ have cofinality μ such that for all $i < \delta$, $j_i < \delta$. Then by continuity $f(M_\delta^1) \leq_{\mathcal{K}} M_\delta^2$. Let C_0 be the set of all such δ . It is easy to check that C_0 is a μ -club. Similarly, let C_1 be the set of all δ such that $f^{-1}(M_\delta^2) \leq_{\mathcal{K}} M_\delta^1$. C_1 is also a μ -club and it is easy to check that $C = C_0 \cap C_1$, and the intersection of two μ -clubs is a μ -club, so the result follows. $\square$

Theorem 6.12. Let μ be a regular cardinal. Assume $\mathcal{K}$ is a μ -AEC, $\lambda = \lambda^{<\mu} \geq \text{LS}(\mathcal{K})$, $I(\lambda, \mathcal{K}) = 1 \leq I(\lambda^+, \mathcal{K}) < 2^{\lambda^+}$. If:

- (1) $\mathcal{K}_\lambda$ has the extension property for δ -chains (see above) for every $\delta < \lambda^+$.
- (2) $\lambda = \lambda^\mu$ and $2^\lambda = \lambda^+$.

Then $\mathcal{K}$ has λ -amalgamation.

Proof. Assume not. By failure of amalgamation and some renaming, we have:

- (*) If $M_1, M_2 \in \mathcal{K}_\lambda$ and $f : M_1 \cong M_2$, there are $M'_l \in \mathcal{K}_\lambda$, $\ell = 1, 2$, with $M_l \leq_{\mathcal{K}} M'_l$, $\text{card}(|M'_l| - |M_l|) = \lambda$, such that there is no $N \in \mathcal{K}_\lambda$ and $g_l : M'_l \rightarrow N$ commuting with f .

In particular (taking $M_1 = M_2$ and f the identity function), the model of size λ is not maximal. By Gregory's theorem (see [Jec03, Theorem 23.2]), the combinatorial principle $\Diamond_{E_\mu}$ holds, where $E_\mu := \{i < \lambda^+ \mid \text{cf } i \geq \mu\}$. With some coding, one can see that $\Diamond_{E_\mu}$ is equivalent to:

- (**) There are $\{\eta_\alpha, \nu_\alpha : \alpha \rightarrow \mid \alpha < \lambda^+\}$, $\{g_\alpha : \alpha \rightarrow \alpha \mid \alpha < \lambda^+\}$ such that for all $\eta, \nu : \lambda^+ \rightarrow 2$, $g : \lambda^+ \rightarrow \lambda^+$, the set $\{\alpha \in E_\mu \mid \eta_\alpha = \eta \restriction \alpha, \nu_\alpha = \nu \restriction \alpha, g_\alpha = g \restriction \alpha\}$ is stationary.

We build a strictly increasing tree $\{M_\eta \mid \eta \in {}^{\leq \lambda^+} 2\}$ such that:

- (1) $|M_\eta| \subseteq \lambda^+$, $\text{card}(M_\eta) = \lambda$, $\ell(\eta) \in |M_{\eta \restriction \ell}|$ for all $\eta \in {}^{< \lambda^+} 2$ and $\ell < 2$.
- (2) If $\eta \in {}^{\leq \lambda^+} 2$ and $\text{cf } \ell(\eta) \geq \mu$, then $M_\eta = \bigcup_{j < \ell(\eta)} M_{\eta \restriction j}$.
- (3) If $|M_{\eta_\delta}| = \delta$, $\eta_\delta \neq \nu_\delta$, and $g_\delta : M_{\eta_\delta} \cong M_{\nu_\delta}$ is an isomorphism, for any $\ell, \ell' < 2$ and any $\nu \supseteq \nu_\delta \restriction \ell'$, g_δ cannot be extended to an embedding of $M_{\eta_\delta \restriction \ell}$ into M_ν .

This is enough. We claim that for any $\eta \neq \nu \in {}^{\lambda^+} 2$, $M_\eta \not\cong M_\nu$. Indeed, assume $f : M_\eta \rightarrow M_\nu$ is an isomorphism. For $i < \lambda^+$, let $f_i := f \restriction M_{\eta \restriction i}$ and let $C := \{i < \lambda^+ \mid f_i : M_{\eta \restriction i} \cong M_{\nu \restriction i}\}$. By Lemma 6.11, C is a μ -club. Also $\{i < \lambda^+ \mid |M_{\eta \restriction i}| = i\}$ is a club so without loss of generality is contained in C . Now the stationary set described by (**) intersects C in unboundedly many places (as it only has points of cofinality μ), hence there is $\delta < \lambda^+$ such that $\eta \restriction \delta \neq \nu \restriction \delta$, $\eta_\delta = \eta \restriction \delta$, $\nu_\delta = \nu \restriction \delta$, $g_\delta = f \restriction \delta$, $\delta = |M_{\eta_\delta}| = |M_{\nu_\delta}|$, and $g_\delta : M_{\eta_\delta} \cong M_{\nu_\delta}$. But f extends g_δ and restricts to an embedding of $M_{\eta \restriction \eta(\delta)}$ into $M_{\nu \restriction \gamma}$, for some $\gamma < \lambda^+$ with $\gamma > \delta$ sufficiently large. This contradicts (3).

This is possible. Take any $M_{<} \in \mathcal{K}$ with $|M_{<}| = \lambda$ for the base case, take unions at limits of cofinality at least μ , and use the extension property for chains (and some renaming) at limits of cofinality less than μ .

Now if one wants to define $M_{\eta \restriction \ell}$ for $\eta \in {}^{\delta 2}$ (assuming by induction that M_ν for all $\nu \in {}^{\leq \delta} 2$ have been defined) take any two strict extensions, unless $|M_\eta| = \delta$, $\eta_\delta \neq \nu_\delta$, $g_\delta : M_{\eta_\delta} \cong M_{\nu_\delta}$ is an isomorphism, and either $\eta = \eta_\delta$, or $\eta = \nu_\delta$. We show what to do when $\eta = \eta_\delta$. The other case is symmetric. Let M'_{η_δ} , M'_{ν_δ} be as described by (*) and let $M_{\eta_\delta \restriction \ell}$, $M_{\nu_\delta \restriction \ell}$ be their appropriate renaming to satisfy (1). Now (*) tells us that (3) is satisfied. $\square$

Remark 6.13. Of course, the set-theoretic hypotheses of Theorem 6.12 can be weakened. For example, it is enough to require $\lambda = \lambda^{<\mu}$ and $\Diamond_{S_\mu}$ or even (by Shelah's more complicated proof) a suitable instance of the weak diamond. It is not clear, however, that it follows from just $2^\lambda < 2^{\lambda^+}$.

REFERENCES

- [AdRo94] Jiří Adámek and Jiří Rosický. **Locally presentable and accessible categories**, London Mathematical Society Lecture Notes Series, **189**. Cambridge UP, Cambridge, 1994, xiv+316 pp.
- [AdRo95] Jiří Adámek and Jiří Rosický, *On preaccessible categories*, *Jour. Pure Appl. Algebra*, **105**, 1995, 225–232.
- [Bal09] John Baldwin, **Categoricity**, University Lecture Series, American Mathematical Society, 2009.
- [BET07] John Baldwin, Paul Eklof, and Jan Trlifaj, $N^\perp$ as an AEC, *Annals Pure Appl. Logic*, **149**, 2007, 25–39.
- [BaFe85] K. Jon Barwise and Solomon Fefferman (Ed.). **Model theoretic logics**, Springer-Verlag 1985.
- [BMK78] John Barwise, Matt Kaufman, and Michael Makkai, *Stationary logic*, *Annals of Mathematical Logic* **13** (1978), 171–224.
- [BMK81] John Barwise, Matt Kaufman, and Michael Makkai, *A correction to ‘Stationary logic’*, *Annals of Mathematical Logic* **20** (1981), 231–232.
- [BBHU08] Itai Ben Yaacov, Alexander Berenstein, C. Ward Henson, and Alexander Usvyatsov, *Model theory for metric structures, Model theory with applications to algebra and analysis*. Vol. 2, London Math. Soc. Lecture Note Ser., vol. 350, Cambridge Univ. Press, Cambridge, 2008, pp. 315–427.
- [BeRo12] Tibor Beke and Jiří Rosický, *Abstract elementary classes and accessible categories*, *Annals Pure Appl. Logic*, **163**, 2012, 2008–2017.
- [Bon14] Will Boney, *Tameness from large cardinals axioms*, *The Journal of Symbolic Logic*, **79**, 2014, 1092–1119.
- [Bon] Will Boney, *A Presentation Theorem for Continuous Logic and Metric Abstract Elementary Classes*, In preparation, URL: <http://arxiv.org/abs/1408.3624>.
- [BoUn] Will Boney and Spencer Unger, *Large cardinal axioms from tameness*, In preparation.
- [BoVaa] Will Boney and Sebastien Vasey, *Chains of Saturated Models in AECs*, Preprint. URL: <http://arxiv.org/abs/1503.08781v3>.
- [BoVab] Will Boney and Sebastien Vasey, *Categoricity in Infinitary Logics*, Preprint. URL: <http://arxiv.org/abs/1508.03316>.
- [BoZa] Will Boney and Pedro Zambrano, *Around the set-theoretic consistency of d -tameness of metric abstract elementary classes*, Preprint. URL: arxiv.org/abs/1508.05529v1.
- [BTRo] Andrew Brooke-Taylor and Jiří Rosický, *Accessible Images Revisited*, arXiv:1506.01986, submitted.
- [Cha68] C. C. Chang. Some remarks on the model theory of infinitary languages, **The Syntax and Semantics of Infinitary Languages**, Springer-Verlag Lecture Notes in Mathematics **72**, edited by J. Barwise (1968) pages 36–63.
- [Dic75] M. A. Dickmann. **Large infinitary logics**. North-Holland Publishing .Co., Amsterdam, 1975, xv+464pp.
- [Gro02] Rami Grossberg, *Classification theory for abstract elementary classes*, *Logic and Algebra* (Yi Zhang, ed.), vol. 302, American Mathematical Society, 2002, pp. 165–204.
- [Gro1X] ———, **A Course in Model Theory**, In Preparation, 201X.
- [GrVa06a] Rami Grossberg and Monica VanDieren, *Galois-stability for tame abstract elementary classes*, *Journal of Mathematical Logic*, **6**, 25–49, 2006.
- [GrVa06b] ———, *Shelah’s categoricity conjecture from a successor for tame abstract elementary classes*, *The Journal of Symbolic Logic*, **71**, 553–568, 2006.

- [GrVa06c] ———, *Categoricity from one successor cardinal in tame abstract elementary classes*, Journal of Mathematical Logic, **6**, 181–201, 2006.
- [GrVa] Rami Grossberg and Sebastien Vasey, *Superstability in abstract elementary classes*, Preprint. URL: <http://arxiv.org/abs/1507.04223v2>.
- [Iwa44] Tsurane Iwamura, *A lemma on directed sets*, Zenkoku Shijo Sugaku Danwakai **262** (1944), 107–111.
- [Jec03] Thomas Jech, **Set Theory**, 3rd ed., Springer, 2003.
- [Law63] F. W. Lawvere, *Functorial semantics of algebraic theories*, Dissertation, Columbia University 1963.
- [Lie11] M. Lieberman, *Category theoretic aspects of abstract elementary classes*, Annals Pure Appl. Logic, **162** (2011), 903–915.
- [LiRo] M. Lieberman and J. Rosický, *Classification theory for accessible categories*, arxiv:1404.2528. To appear in J. of Symb. Logic.
- [LiRoa] ———, *Metric AECs as accessible categories*, submitted, arXiv:1504.02660.
- [MV] Menachem Magidor and Jouko Väänänen, *On Löwenheim-Skolem-Tarski numbers for extensions of first order logic*, Jour. Math. Logic **11** (2011), 87–113.
- [MaPa89] M. Makkai and R. Paré, *Accessible Categories: The Foundations of Categorical Model Theory*, AMS 1989.
- [Mac71] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed., Graduate Texts in Mathematics **5**, Springer 1998.
- [Ros81] J. Rosický, *Concrete categories and infinitary languages*, Journal of Pure and Applied Algebra, **22** (1981), 309–339.
- [Ros97] ———, *Accessible categories, saturation and categoricity*, Journal of Symb. Logic, **62**, 891–901, 1997.
- [Sh43] Saharon Shelah, *Generalized quantifiers and compact logics*, Transactions of the American Mathematical Society **204** (1975), 342–364.
- [Sh:h] Saharon Shelah, **Classification Theory for Abstract Elementary Classes**, vol. 1 & 2, Mathematical Logic and Foundations, no. 18 & 20, College Publications, 2009.
- [Sh87a] Saharon Shelah Classification theory for nonelementary classes I, The number of uncountable models of $\psi \in L_{\omega_1, \omega}$. Part A, *Israel Journal of Mathematics*, **46**:212–240, 1983.
- [Sh87b] Saharon Shelah. Classification theory for nonelementary classes. I. The number of uncountable models of $\psi \in L_{\omega_1, \omega}$. Part B, *Israel Journal of Mathematics*, **46**:241–273, 1983.
- [Sh88] Saharon Shelah, Classification of nonelementary classes II, Abstract elementary classes, In *Classification theory (Chicago, IL, 1985)*, volume 1292 of *Lecture Notes in Mathematics*, pages 419–497. Springer, Berlin, 1987. Proceedings of the USA–Israel Conference on Classification Theory, Chicago, December 1985; ed. Baldwin, J.T.
- [Sh394] Saharon Shelah, *Categoricity for abstract classes with amalgamation*, Annals of Pure and Applied Logic, **98**, 261–294, 1999.
- [Sh702] Saharon Shelah, *On What I Do Not Understand (and have something to say), model theory*, Mathematica Japonica **51** (2000), 329–377.
- [Sh1019] Saharon Shelah, *Model theory for a compact cardinal*, Preprint. URL: <http://arxiv.org/abs/1303.5247v2>.
- [Vasa] Sebastien Vasey, *Independence in Abstract Elementary Classes*, Preprint. URL: <http://arxiv.org/abs/1503.01366v5>.
- [Vasb] ———, *Infinitary stability theory*, Preprint. URL <http://arxiv.org/abs/1412.3313v5>.

[Vasc] Sebastien Vasey, *Forking and superstability in tame AECs*, The Journal of Symbolic Logic, To appear. URL: <http://arxiv.org/abs/1405.7443v3>.

E-mail address: wboney@math.harvard.edu

URL: <http://math.harvard.edu/~wboney/>

MATHEMATICS DEPARTMENT, HARVARD UNIVERSITY, CAMBRIDGE, MASSACHUSETTS, USA

E-mail address: rami@cmu.edu

URL: <http://math.cmu.edu/~rami/>

DEPARTMENT OF MATHEMATICAL SCIENCES, CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PENNSYLVANIA, USA

E-mail address: lieberman@math.muni.cz

URL: <http://www.math.muni.cz/~lieberman/>

DEPARTMENT OF MATHEMATICS AND STATISTICS, FACULTY OF SCIENCE, MASARYK UNIVERSITY, BRNO, CZECH REPUBLIC

E-mail address: rosicky@math.muni.cz

URL: <http://www.math.muni.cz/~rosicky/>

DEPARTMENT OF MATHEMATICS AND STATISTICS, FACULTY OF SCIENCE, MASARYK UNIVERSITY, BRNO, CZECH REPUBLIC

E-mail address: sebv@cmu.edu

URL: <http://math.cmu.edu/~svasey/>

DEPARTMENT OF MATHEMATICAL SCIENCES, CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PENNSYLVANIA, USA