The Formal Specification Column

BY

Hartmut Ehrig

Technical University of Berlin, Department of Computer Science
Franklinstraße 28/29, D-10587 Berlin, Germany
ehrig@cs.tu-berlin.de

Some Nuances of Many-sorted Universal Algebra: A Review*

Andrzej Tarlecki

Institute of Informatics, University of Warsaw
and
Institute of Computer Science, Polish Academy of Sciences
Warsaw, Poland.
tarlecki@mimuw.edu.pl

Abstract

It has been a common belief that the standard results of universal algebra as developed since the work of Birkhoff and others in the thirties carry over without much change to the framework of many-sorted algebras. Perhaps the only exception widely noticed by the community is the care needed in the treatment of many-sorted equational logic. However, while the standard results remain valid in essence in the many-sorted frameworks, some nuances and technicalities require considerably more care in formulation and

*This work was funded in part by by the Polish Ministry of Science and Higher Education, grant N206 493138, and by the European IST FET programme under the IST-2005-016004 SENSORIA project.

proof of the results. We give some examples of this, indicating how equational calculus, Birkhoff's variety theorem and interpolation results should be adjusted for many-sorted algebras.

1 Introduction

Since the work of Birkhoff and others in the thirties, universal algebra has been developed as a highly influential branch of mathematics. We refer to [9] for a handbook presentation of many ideas, concepts and results of universal algebra. Most of the material there, as in most other mathematical literature on universal algebra, is presented in the context of single-sorted (homogeneous) algebras. However, when the theory of algebraic specifications has been developed as a branch of computer science (see for instance [10, 6, 12, 17] for relatively early presentations of this area) it became clear that a version of universal algebra based on many-sorted (heterogeneous) algebras is needed, as perhaps first explicitly introduced in [11, 2].

It has been widely accepted then that many-sorted universal algebra does not differ much from its single-sorted version. It came as a surprise when it was realised that even the basic equational calculus needs to be adjusted when used in the many-sorted framework. An elegant overall solution came in [7], with a number of variants developed independently by other authors, see for instance [6], and heated discussions over various details at the time (see e.g. [5, 8]). Later on, other differences between many-sorted and single-sorted algebraic frameworks have been pointed at.

In this note we try to indicate some of the standard results and ideas of universal algebra which do not carry over to the many-sorted framework in an entirely obvious manner. We start though by recalling in Sect. 2 the basic definitions and few basic facts of universal algebra in many-sorted formulation. Even though this material is quite standard, we present it not only to introduce notation but also to point out already here at some concepts that should be considered with more care in the many-sorted framework. Then, in Sect. 3, we have a look at the famous Birkhoff's variety theorem, characterising equationally definable classes as varieties [1]. It turns out that although the formulation of the theorem carries over to the many-sorted framework, significant technical difficulties are caused by nuances concerned with infinite sets of sorts and algebras that have carriers of some sorts empty. In Sect. 4 we point at differences between single-sorted and many-sorted equational calculus, following largely [7]. Finally, in Sect. 5, we look at interpolation properties, where again, many-sorted framework requires different than standard formulations of the interpolation properties for first-order [3] and for equational logic [14].

We do not claim any technical novelty here: the results presented are either available in the literature (we apologise for the highly incomplete list of references) or in the folklore of the field. As mentioned above, the issues concerning many-sorted equational calculus were resolved by the mid-eighties; careful formulations of Birkhoff's variety theorem in many-sorted context were given about the same time (e.g. in [6]); a more complete understanding of interpolation for many-sorted logic came somewhat later (see [13, 3]). However, we think that exposing the nuances of many-sorted universal algebra should be useful to clarify many misunderstandings and common inaccuracies that still may be found in the literature.

Acknowledgements: Many thanks to Bartek Klin and Jiři Adámek for alerting me to the problems with precise formulation of Birkhoff's variety theorem in the many-sorted framework, and to Tomasz Brendos and Jarosław Tworek for early discussions on this topic.

Very special thanks to Don Sannella for long-term collaboration on these and related topics. Some parts of the material presented here are included or hinted at in our forthcoming monograph [15].

2 Basic Definitions and Facts

2.1 Many-sorted Sets

Let S be any set; we think of elements of S as *sort names*, or *sorts* for short.

An S*-sorted set* is an S-indexed family of sets $X = \langle X_s \rangle_{s \in S}$. We say that such an S-sorted set X is *empty* if X_s is empty for all $s \in S$. The empty S-sorted set will be written (ambiguously) as $\emptyset$. We say that X is *everywhere non-empty* if $X_s \neq \emptyset$ for all $s \in S$; otherwise we say that X is *somewhere empty*. Clearly, if S has at least two elements, there are S-sorted sets that are neither empty nor everywhere non-empty.

S-sorted set X is *finite* if X_s is finite for all $s \in S$ and $X_s = \emptyset$ for almost all $s \in S$ (that is, for all but a finite number of $s \in S$, $X_s = \emptyset$).

Let $X = \langle X_s \rangle_{s \in S}$ and $Y = \langle Y_s \rangle_{s \in S}$ be S-sorted sets. Union, intersection, Cartesian product, disjoint union, inclusion (subset) and equality of X and Y are defined component-wise in the obvious manner.

An S*-sorted function* $f : X \to Y$ is an S-indexed family of functions $f = \langle f_s : X_s \to Y_s \rangle_{s \in S}$; X is called the *domain* (or *source*) of f, and Y is called its *codomain* (or *target*). An S-sorted function $f : X \to Y$ is an *identity* (*inclusion*, *surjection*, *injection*, *bijection*, etc) if for every $s \in S$, the function $f_s : X_s \to Y_s$ is an identity (inclusion, surjection, injection, bijection, etc). The identity S-sorted function on X will be written as $id_X : X \to X$.

S-sorted functions compose in the usual, component-wise manner; we write their composition in the diagrammatic order, using ";" (semicolon): for $f : X \to Y$

and $g\colon Y \to Z$, their composition is $f;g\colon X \to Z$. The category of S-sorted sets with S-sorted functions as morphisms will be written as $\mathbf{Set}^S$.

All other standard concepts of set theory (in particular, *image* and *coimage* of a set w.r.t. a function, *binary relation*, *equivalence*, *quotient set*, etc) carry over to S-sorted context as above, in the expected, component-wise manner, with the standard notations retained. For instance, given an S-sorted set X and an S-sorted equivalence $\equiv \subseteq X \times X$, we write $x \in X$ instead of $x \in X_s$ for some sort $s \in S$, and then $[x]_\equiv$ for $[x]_{\equiv_s}$, the equivalence class of x w.r.t. $\equiv_s$.

2.2 Many-sorted Algebras

A *(many-sorted) signature* is a pair $\Sigma = \langle S, \Omega \rangle$, where:

- S is a set (of sort names); and
- Ω is an $(S^* \times S)$-sorted set (of operation names)

where S^* is the set of finite (including empty) sequences of elements of S. When the signature is understood, we write $f\colon s_1 \times \cdots \times s_n \to s$ whenever $s_1, \ldots, s_n, s \in S$ and $f \in \Omega_{\langle s_1 \cdots s_n, s \rangle}$.

To avoid minor technical complications, we will assume that in each signature, the sets in Ω are mutually disjoint (so that no overloading of operation names is allowed). This is not an essential assumption for any of the technical developments below.

For the rest of this section, let $\Sigma = \langle S, \Omega \rangle$ be a signature.

A Σ-*algebra* A consists of:

- an S-sorted set $|A|$ of *carrier sets* (or *carriers*); and
- for $f\colon s_1 \times \cdots \times s_n \to s$ in Σ, a function (or *operation*) $f_A\colon |A|_{s_1} \times \cdots \times |A|_{s_n} \to |A|_s$.

We say that A is *everywhere non-empty* if $|A|$ is everywhere non-empty, i.e., $|A|_s \neq \emptyset$ for all sorts $s \in S$; otherwise A is *somewhere empty*.

Let A and B be Σ-algebras. B is a *subalgebra* of A if:

- $|B| \subseteq |A|$; and
- for $f\colon s_1 \times \cdots \times s_n \to s$ in Σ and $b_1 \in |B|_{s_1}, \ldots, b_n \in |B|_{s_n}$, $f_B(b_1, \ldots, b_n) = f_A(b_1, \ldots, b_n)$.

In other words, any subalgebra B of A is given by an S-sorted subset $|B|$ of $|A|$ that is closed under Σ-operations as defined in A. Subalgebras of any algebra A, ordered by inclusion of their carriers, form a complete lattice. In particular, given any set $X \subseteq |A|$, there is the least subalgebra of A that contains X; we write it as $\langle X \rangle_A$ and call it the *subalgebra of A generated by X*. This is well-defined since the intersection of any family of subalgebras of A is a subalgebra of A, which is not true when everywhere non-empty algebras are considered: intersection of everywhere non-empty subalgebras of A may be somewhere empty.

Given a family $\langle A_i \rangle_{i \in I}$ of Σ-algebras (indexed by any set I of indexes), its *product* $\prod_{i \in I} A_i$ is a Σ-algebra P defined as follows:

- $|P| = \prod_{i \in I} |A_i|$; and
- for $f : s_1 \times \cdots \times s_n \to s$ in Σ and $f_1 \in |P|_{s_1}, \ldots, f_n \in |P|_{s_n}$, $f_P(f_1, \ldots, f_n)(i) = f_{A_i}(f_1(i), \ldots, f_n(i))$ for all $i \in I$.

We use the usual notation $A \times B$ for binary products. The product of the empty family of Σ-algebras (the family indexed by $I = \emptyset$) is the "singleton" algebra, which has a singleton carrier set for each sort; we write it as $\mathbf{1}_\Sigma$ (of course, this determines $\mathbf{1}_\Sigma$ up to an isomorphism only, see below).

A Σ-*homomorphism* $h : A \to B$ between Σ-algebras A and B is an S-sorted function $h : |A| \to |B|$ which preserves the operations of Σ, i.e. such that for all $f : s_1 \times \cdots \times s_n \to s$ in Σ and $a_1 \in |A|_{s_1}, \ldots, a_n \in |A|_{s_n}$, $h_s(f_A(a_1, \ldots, a_n)) = f_B(h_{s_1}(a_1), \ldots, h_{s_n}(a_n))$.

Composition of Σ-homomorphisms (as functions) is a Σ-homomorphism and the identities are Σ-homomorphisms. The category of Σ-algebras with Σ-homomorphisms as morphisms will be denoted by $\mathbf{Alg}(\Sigma)$. $\mathbf{Alg}_{NE}(\Sigma)$ stands for its full subcategory of everywhere non-empty algebras. A Σ-homomorphism is a Σ-*isomorphism* if it is bijective (these are exactly isomorphisms in $\mathbf{Alg}(\Sigma)$).

If $h : A \to B$ is a Σ-homomorphism then for any subalgebra A' of A, the image of A' under h, written $h(A')$, is a subalgebra of B, and for any subalgebra B' of B, the coimage of B' w.r.t. h, written $h^{-1}(B')$ is a subalgebra of A. The latter property does not carry over to everywhere non-empty algebras: the image of an everywhere non-empty subalgebra is everywhere non-empty, but the coimage of an everywhere non-empty subalgebra may be somewhere empty.

A Σ-*congruence* on a Σ-algebra A is an (S-sorted) equivalence $\equiv$ on $|A|$ which respects the operations of Σ: for all $f : s_1 \times \cdots \times s_n \to s$ in Σ and $a_1, a'_1 \in |A|_{s_1}$, $\ldots$, $a_n, a'_n \in |A|_{s_n}$, if $a_1 \equiv_{s_1} a'_1$ and $\ldots$ and $a_n \equiv_{s_n} a'_n$ then $f_A(a_1, \ldots, a_n) \equiv_s f_A(a'_1, \ldots, a'_n)$.

Given a Σ-congruence $\equiv$ on a Σ-algebra A, the *quotient algebra* $A/\!\equiv$ of A modulo $\equiv$ is the Σ-algebra Q defined by:

- $|Q| = |A|/\!\equiv$; and
- for $f : s_1 \times \cdots \times s_n \to s$ and $a_1 \in |A|_{s_1}, \ldots, a_n \in |A|_{s_n}$, $f_Q([a_1]_\equiv, \ldots, [a_n]_\equiv) = [f_A(a_1, \ldots, a_n)]_\equiv$.

Given a Σ-homomorphism $h : A \to B$, its kernel $ker(h) = \{\langle a, a' \rangle \mid h(a) = h(a')\} \subseteq |A| \times |A|$ is a Σ-congruence. The quotient algebra $A/ker(h)$ is isomorphic to $h(A)$ via the isomorphism that maps any equivalence class $[a]_{ker(h)}$, $a \in |A|$, to $h(a)$.

A non-empty family $\langle A_i \rangle_{i \in I}$ of Σ-algebras is *directed* if for all $i_1, i_2 \in I$ there is $i \in I$ such that A_{i_1} and A_{i_2} are subalgebras of A_i. Given such a family, its *(directed) sum* $\coprod_{i \in I} A_i$ is the Σ-algebra U defined as follows:

- $|U| = \bigcup_{i \in I} |A_i|$, and
- for $f \colon s_1 \times \cdots \times s_n \to s$ in Σ, $f_U = \bigcup_{i \in I} f_{A_i}$.

2.3 Equations

Let $\Sigma = \langle S, \Omega \rangle$ be an algebraic signature, fixed throughout this section.

Given an S-sorted set X such that for all distinct $s, s' \in S$, X_s and $X_{s'}$ are disjoint[1], Σ-*terms* with variables in X are defined as usual, and so is the usual Σ-algebra $T_\Sigma(X)$ of Σ-terms with variables in X. Let $\eta_X \colon X \to |T_\Sigma(X)|$ be the inclusion. Given a term $t \in |T_\Sigma(X)|$, we write $FV(t) \subseteq X$ for the (finite) S-sorted set of variables that occur in t; clearly $FV(t) \subseteq X$, but in general the inclusion may be proper. We say that t is *ground* if $FV(t) = \emptyset$. The algebra of ground terms is written as T_Σ.

If X is everywhere non-empty then so is the term algebra $T_\Sigma(X)$. In general though. term algebras $T_\Sigma(X)$ may be somewhere empty. All Σ-algebras are everywhere non-empty (so that $\mathbf{Alg}(\Sigma)$ and $\mathbf{Alg}_{NE}(\Sigma)$ coincide) if and only if the algebra T_Σ of ground terms is everywhere non-empty.

We say that a set of sorts $S_0 \subseteq S$ *makes a sort* $s \in S$ *non-void in* Σ if there is a Σ-term of sort s with variables of sorts in S_0 only. A set of sorts $S' \subseteq S$ is *non-void in* Σ if the empty set of sorts makes each $s \in S'$ non-void in Σ (that is, there is a ground Σ-term of sort s). A set of sorts $S' \subseteq S$ is *almost non-void in* Σ if there is a finite $S_0 \subseteq S'$ such that S_0 makes each sort $s \in S'$ non-void in Σ. The term algebra $T_\Sigma(X)$ is everywhere non-empty if the set of sorts $\{s \in S \mid X_s \neq \emptyset\}$ makes each sort $s \in S$ non-void in Σ. The set of all sorts in Σ is almost non-void in Σ if and only if there exists a finite S-sorted set X such that $T_\Sigma(X)$ is everywhere non-empty.

The algebra $T_\Sigma(X)$ if *free in* $\mathbf{Alg}(\Sigma)$ *with generators* X, that is, for any Σ-algebra A and function (*valuation* of variables) $v \colon X \to |A|$, there exists a unique Σ-homomorphism $v^\sharp \colon T_\Sigma(X) \to A$ that extends v, that is, such that $\eta_X;v^\sharp = v$. Then, given a term $t \in |T_\Sigma(X)|$, we write $t_A(v) = v^\sharp(t)$ for the *value* of t in A under v. The value of a term depends only on the valuation of the variables that occur free in it: given a term $t \in |T_\Sigma(X)|$, Σ-algebra A, and two valuations $v_1, v_2 \colon X \to |A|$ such that v_1 and v_2 coincide on $FV(t)$, we have $t_A(v_1) = t_A(v_2)$.

[1]This is assumed without mention in the following as well, whenever an S-sorted set is used as a set of variables. We add this minor assumption to avoid some technical troubles below, following a similar assumption about the operation names in the signatures considered. For notational convenience, we will also assume that variables are distinct from constants in the signature.

In particular, given two S-sorted sets X and Y, and a *substitution* $\theta\colon X \to |T_\Sigma(Y)|$ of Σ-terms (with variables Y) for variables in X, we have the unique homomorphism $\theta^\sharp\colon T_\Sigma(X) \to T_\Sigma(Y)$, which substitutes terms $\theta(x)$ for variables $x \in X$ in any Σ-term $t \in |T_\Sigma(X)|$; we write $t[\theta]$ for $\theta^\sharp(t) = t_{T_\Sigma(Y)}(\theta)$.

A Σ-*equation* $\forall X.\, t = t'$ consists of:

- an S-sorted set X (of variables), and
- two Σ-terms $t, t' \in |T_\Sigma(X)|_s$ for some sort $s \in S$.

An equation $\forall X.\, t = t'$ is *finitary* if X is finite; it is *ground* if $X = \emptyset$.

A *naive* Σ-equation $t = t'$ consists of two Σ-terms $t, t' \in |T_\Sigma(X)|_s$ for some S-sorted set X (of variables) and sort $s \in S$; we identify naive Σ-equation $t = t'$ with the Σ-equation $\forall X.\, t = t'$, where $X = FV(t) \cup FV(t')$ is the set of variables that actually occur in terms t, t'. Clearly, naive equations are finitary.

A Σ-algebra A *satisfies* (or, *is a model of*) a Σ-equation $\forall X.\, t = t'$, written $A \models_\Sigma \forall X.\, t = t'$, if for every valuation $v\colon X \to |A|$, $t_A(v) = t'_A(v)$.

As usual, we omit the subscript Σ whenever convenient, and for any class $\mathcal{A}$ of Σ-algebras, Σ-algebra A, set Φ of Σ-equations and Σ-equation φ, we write $\mathcal{A} \models \varphi$, $A \models \Phi$ and $\mathcal{A} \models \Phi$ with the obvious (conjunctive) meaning.

More crucially, we write $\Phi \models \varphi$ whenever φ is a *semantic consequence* of Φ, that is, for all Σ-algebras $A \in |\mathbf{Alg}(\Sigma)|$, if $A \models \Phi$ then also $A \models \varphi$.

Restricting attention to everywhere non-empty algebras, we write $\Phi \models^{NE} \varphi$ whenever φ is a *semantic consequence* of Φ for everywhere non-empty algebras, that is, for all everywhere non-empty Σ-algebras $A \in |\mathbf{Alg}_{NE}(\Sigma)|$, if $A \models \Phi$ then also $A \models \varphi$. Clearly, the semantic entailment implies semantic entailment for everywhere non-empty algebras: if $\Phi \models \varphi$ then $\Phi \models^{NE} \varphi$ — but the opposite implication fails in general.

Given a class $\mathcal{A}$ of Σ-algebras, we write $ETh(\mathcal{A})$ for the set[2] of all Σ-equations that are satisfied in all algebras in $\mathcal{A}$; $ETh_{fin}(\mathcal{A})$ stands for the set of all finitary equations that hold in all algebras in $\mathcal{A}$, and $ETh_{naive}(\mathcal{A})$ for the set of all naive equations that hold in all algebras in $\mathcal{A}$. We have $ETh_{naive}(\mathcal{A}) \subseteq ETh_{fin}(\mathcal{A}) \subseteq ETh(\mathcal{A})$, and both inclusions are proper (unless the signature is empty).

Given a set Φ of Σ-equations, we write $Mod(\Phi)$ for the class of all algebras that satisfy Φ; classes of this form are called *equationally definable*. So, $ETh(Mod(\Phi))$ is the set of all semantic consequences of Φ.

We also write $Mod_{NE}(\Phi)$ for the class of all everywhere non-empty algebras that satisfy Φ, and call such classes *everywhere non-empty equationally definable*. In general, everywhere non-empty equationally definable class need not be equationally definable, but an equationally definable class that consists of everywhere

[2]This indeed is a set if we presume that variables in equations are selected from a predefined, sufficiently large vocabulary.

non-empty algebras only is everywhere non-empty definable. Everywhere non-empty equationally definable classes are everywhere non-empty definable by sets of naive (and hence finitary) equations.

Given a Σ-equation $\forall X.\, t = t'$, if a Σ-algebra A satisfies the naive Σ-equation $t = t'$ then it also satisfies $\forall X.\, t = t'$; that is: $t = t' \models \forall X.\, t = t'$. The opposite fails in general though. However, if A is everywhere non-empty then $A \models_\Sigma \forall X.\, t = t'$ iff $A \models_\Sigma t = t'$. Consequently, $\forall X.\, t = t'$ and $t = t'$ are semantically equivalent for everywhere non-empty algebras, so that if only everywhere non-empty algebras are considered, any equation may be replaced by a naive (and hence finitary) one. For arbitrary algebras, this is not always possible:

Fact 2.1. *Consider a Σ-equation $\forall X.\, t = t'$. If the set $\{s \in S \mid X_s \neq \emptyset\}$ is almost non-void in Σ then there is a finite $X' \subseteq X$ such that the finitary Σ-equation $\forall X'.\, t = t'$ is semantically equivalent to $\forall X.\, t = t'$.*

Proof sketch. *Let $S' = \{s \in S \mid X_s \neq \emptyset\}$, and let $S_0 \subseteq S'$ be a finite set of sorts that makes each sort $s \in S'$ non-void in Σ. Then take X' to be $FV(t) \cup FV(t')$ with one variable from X_{s_0} added for each sort $s_0 \in S_0$.* □

2.4 Signature morphisms

A *signature morphism* $\sigma\colon \Sigma \to \Sigma'$ between signatures $\Sigma = \langle S, \Omega\rangle$ and $\Sigma' = \langle S', \Omega'\rangle$ is a pair $\sigma = \langle \sigma_{sorts}, \sigma_{opns}\rangle$ where $\sigma_{sorts}\colon S \to S'$ and σ_{opns} is a family of functions that map Σ-operation names to Σ'-operation names respecting their arities and result sorts, that is $\sigma_{opns} = \langle \sigma_{w,s}\colon \Omega_{w,s} \to \Omega'_{\sigma^*_{sorts}(w),\sigma_{sorts}(s)}\rangle_{w\in S^*, s\in S}$ (where for $w = s_1 \ldots s_n \in S^*$, $\sigma^*_{sorts}(w) = \sigma_{sorts}(s_1)\ldots\sigma_{sorts}(s_n)$). A signature morphism $\sigma\colon \Sigma \to \Sigma'$ is a *signature inclusion* $\sigma\colon \Sigma \hookrightarrow \Sigma'$ if σ_{sorts} is an inclusion and $\sigma_{w,s}$ is an inclusion for all $w \in S^*, s \in S$. When no confusion may arise, we omit all the subscripts and write σ for all the components of the morphism σ.

It is easy to check that signature morphisms compose in the natural way. This yields the category **AlgSig** of algebraic many-sorted signatures and their morphisms. This category is well known to be cocomplete, with the empty signature as the initial object, coproducts defined as disjoint unions of signatures, and coequalisers given by identifying sort and operation names in a minimal way to coequalise the morphisms involved. Consequently, any pushout in **AlgSig**

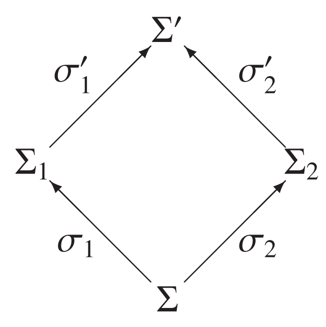

is constructed so that Σ' is the disjoint union of Σ_1 and Σ_2 but with the sorts and operation names that have a common source in Σ identified.

Signature morphism $\sigma\colon \Sigma \to \Sigma'$ determines a *σ-reduct functor* $\_|_\sigma\colon \mathbf{Alg}(\Sigma') \to \mathbf{Alg}(\Sigma)$ defined as follows:

- For any Σ'-algebra $A' \in |\mathbf{Alg}(\Sigma')|$, $A'|_\sigma \in |\mathbf{Alg}(\Sigma)|$ is defined by:
  - $|A'|_\sigma|_s = |A'|_{\sigma(s)}$ for all $s \in S$; and
  - for all $f\colon s_1 \times \cdots \times s_n \to s$ in Σ, $f_{A'|_\sigma} = \sigma(f)_{A'}$.
- For any Σ'-homomorphism $h'\colon A' \to B'$, $h'|_\sigma = \langle h'_{\sigma(s)}\rangle_{s\in S} : A'|_\sigma \to B'|_\sigma$.

In general, the reduct functor $\_|_\sigma\colon \mathbf{Alg}(\Sigma') \to \mathbf{Alg}(\Sigma)$ need not be surjective (on objects) even if σ is injective. However, given an injective signature morphism σ, the reduct w.r.t. σ is surjective on everywhere non-empty algebras, so that we have a functor $\_|_\sigma\colon \mathbf{Alg}_{NE}(\Sigma') \to \mathbf{Alg}_{NE}(\Sigma)$ which is surjective on objects. In fact, such reduct functors are also surjective on morphisms between everywhere non-empty algebras. Moreover, the following property holds:

Fact 2.2. *Consider an injective signature morphism $\sigma\colon \Sigma \to \Sigma'$.*

For any everywhere non-empty Σ'-algebra $A' \in |\mathbf{Alg}_{NE}(\Sigma')|$ and Σ-algebra B, if $A'|_\sigma$ is a subalgebra of B then there exists an everywhere non-empty Σ'-algebra B' such that $B'|_\sigma = B$ and A' is a subalgebra of B'.

For any (everywhere non-empty) Σ'-algebra $A' \in |\mathbf{Alg}(\Sigma')|$, Σ-algebra B and surjective Σ-homomorphism $h\colon B \to A'|_\sigma$, there exist an (everywhere non-empty) Σ'-algebra B' and a surjective Σ'-homomorphism $h'\colon B' \to A'$ such that $h'|_\sigma = h$ (and $B'|_\sigma = B$). □

A (constructive) proof may be extracted from [14]; for the first part, the restriction to everywhere non-empty algebras cannot be omitted since the result may fail if we allow A' considered there to be somewhere empty.

The assignments $\Sigma \mapsto \mathbf{Alg}(\Sigma)$ and $\sigma \mapsto \_|_\sigma$ yield a (contravariant) functor from the category of signatures **AlgSig** to the (quasi-)category **Cat** of all categories, $\mathbf{Alg}\colon \mathbf{AlgSig}^{op} \to \mathbf{Cat}$. Most crucially, this functor is continuous (maps colimits of signatures to limits of categories), which in particular means that the following *amalgamation property* holds.

Any pushout in the category **AlgSig** of signatures

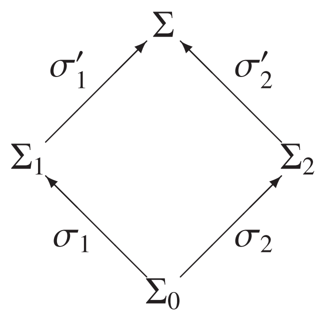

admits amalgamation, that is:

- for any algebras $A_1 \in |\mathbf{Alg}(\Sigma_1)|$ and $A_2 \in |\mathbf{Alg}(\Sigma_2)|$ such that $A_1\big|_{\sigma_1} = A_2\big|_{\sigma_2}$, there exists a unique algebra $A \in |\mathbf{Alg}(\Sigma)|$ such that $A\big|_{\sigma'_1} = A_1$ and $A\big|_{\sigma'_2} = A_2$; and
- for any homomorphisms $h_1 \colon A_{11} \to A_{12}$ in $\mathbf{Alg}(\Sigma_1)$ and $h_2 \colon A_{21} \to A_{22}$ in $\mathbf{Alg}(\Sigma_2)$ such that $h_1\big|_{\sigma_1} = h_2\big|_{\sigma_2}$, there exists a unique Σ-homomorphism $h \colon A_1 \to A_2$ in $\mathbf{Alg}(\Sigma)$ such that $h\big|_{\sigma'_1} = h_1$ and $h\big|_{\sigma'_2} = h_2$.

Clearly, amalgamation property also holds when only everywhere non-empty algebras are considered.

Any signature morphism $\sigma \colon \Sigma \to \Sigma'$ determines translation of Σ-terms to Σ'-terms. Namely, given $t \in |T_\Sigma(X)|_s$, $\sigma(t) \in |T_{\Sigma'}(\sigma(X))|_{\sigma(s)}$, where for $s' \in S'$, $(\sigma(X))_{s'} = \bigcup\{X_s \mid s \in S, f(s) = s'\}$, and $\sigma(t)$ is obtained from t by substituting $\sigma(f)$ for f, for all operation names f. This naturally extends to Σ-equations: $\sigma(\forall X.\, t_1 = t_2) = \forall\sigma(X).\, \sigma(t_1) = \sigma(t_2)$.

Translation of equations and reducts of algebras w.r.t. a signature morphism are closely linked with each other by so-called *satisfaction property*. For any signature morphism $\sigma \colon \Sigma \to \Sigma'$, Σ'-algebra A' and Σ-equation φ:

$$A'\big|_\sigma \models \varphi \text{ iff } A' \models \sigma(\varphi)$$

In particular, this implies that semantic consequence is preserved under translation w.r.t. signature morphisms: given a set Φ of Σ-equations and Σ-equation φ, for any signature morphism $\sigma \colon \Sigma \to \Sigma'$, if $\Phi \models \varphi$ then $\sigma(\Phi) \models \sigma(\varphi)$. The opposite implication may fail though. However, if the reduct functor $\_\big|_\sigma \colon \mathbf{Alg}(\Sigma') \to \mathbf{Alg}(\Sigma)$ is surjective on objects then also the opposite implication holds and we have equivalence:

$$\Phi \models \varphi \text{ iff } \sigma(\Phi) \models \sigma(\varphi)$$

Consequently, if we restrict attention to everywhere non-empty algebras and injective signature morphisms σ, so that reduct functors are surjective on objects, then the equivalence holds:

$$\Phi \models^{NE} \varphi \text{ iff } \sigma(\Phi) \models^{NE} \sigma(\varphi)$$

3 Birkhoff's Variety Theorem

Let $\Sigma = \langle S, \Omega\rangle$ be an algebraic signature.

A class of Σ-algebras $\mathcal{V} \subseteq |\mathbf{Alg}(\Sigma)|$ is a *variety* if it is closed under:

- products: if $A_i \in \mathcal{V}$ for $i \in I$, then $\prod_{i\in I} A_i \in \mathcal{V}$,
- subalgebras: if $A \in \mathcal{V}$ and A_0 is a subalgebra of A then $A_0 \in \mathcal{V}$, and

- homomorphic images: if $A \in \mathcal{V}$ and $h\colon A \to B$ is a Σ-homomorphism then $h(A) \in \mathcal{V}$.

Varieties are closed under isomorphism: if $A \in \mathcal{V}$ and B is isomorphic to A then $B \in \mathcal{V}$ (since it is the image of A under the isomorphism between them). Moreover, every variety $\mathcal{V}$ is non-empty: it contains the product of the empty family of algebras, $\mathbf{1}_\Sigma \in \mathcal{V}$ (Σ-algebra with a one-element carrier of each sort).

Fact 3.1. *Given a class of Σ-algebras $\mathcal{A} \subseteq |\mathbf{Alg}(\Sigma)|$, there is the least Σ-variety that contains $\mathcal{A}$, which consists of all homomorphic images of subalgebras of products of algebras in $\mathcal{A}$.* □

A class $\mathcal{V}$ of Σ-algebras is *finitary* if it is closed under directed sums: if $\langle A_i\rangle_{i\in I}$ is a directed family of algebras in $\mathcal{V}$ then also $\coprod_{i\in I} A_i \in \mathcal{V}$.

A class of everywhere non-empty Σ-algebras $\mathcal{V} \subseteq |\mathbf{Alg}_{NE}(\Sigma)|$ is an *everywhere non-empty variety* if it is closed under products, homomorphic images and everywhere non-empty subalgebras: if $A \in \mathcal{V}$ and A' is an everywhere non-empty subalgebra of A then $A' \in \mathcal{V}$. Note that an everywhere non-empty variety need not be a variety in general, but if a variety consists of everywhere non-empty algebras only (this is the case if the algebra of ground Σ-terms is everywhere non-empty) then it is an everywhere non-empty variety.

It is easy to check that the class of models of any Σ-equation is closed under products, subalgebras and homomorphic images. Moreover, the class of models of any finitary equation is closed under directed sums. This yields:

Lemma 3.2. *For any set Φ of Σ-equations, $Mod(\Phi)$ is a variety. Moreover, if all equations in Φ are finitary then $Mod(\Phi)$ is finitary as well.* □

Corollary 3.3. *For any set Φ of Σ-equations, $Mod_{NE}(\Phi)$ is a finitary everywhere non-empty variety.*
Proof sketch. *For everywhere non-empty algebras, any equation is equivalent to a naive (and hence finitary) equation.* □

We aim now at a complete characterisation of equationally definable classes as varieties.

Lemma 3.4. *Let $\mathcal{V}$ be a variety of Σ-algebras and X be an S-sorted set. Then there exists a Σ-algebra $F^{\mathcal{V}}(X) \in \mathcal{V}$ with a function $\eta^{\mathcal{V}}_X\colon X \to |F^{\mathcal{V}}(X)|$ that is free over X in $\mathcal{V}$, that is, such that for any algebra $A \in \mathcal{V}$ and function $v\colon X \to |A|$, there exists a unique Σ-homomorphism $v^\sharp\colon F^{\mathcal{V}}(X) \to A$ with $\eta^{\mathcal{V}}_X;v^\sharp = v$.*
Proof sketch. *Consider a set of Σ-algebras $A_i \in \mathcal{V}$, $i \in I$, with functions $v_i\colon X \to |A_i|$ such that for any algebra $A \in \mathcal{V}$ and function $v\colon X \to A$, for some $i \in I$, we have an injective Σ-homomorphism $h\colon A_i \to A$ such that $v_i;h = v$. (Such a set exists: the algebras A_i may be chosen for instance as quotients of the term algebra*

$T_\Sigma(X)$.) Let then $P = \prod_{i\in I} A_i$, and $\eta^{\mathcal{V}}_X \colon X \to |P|$ be given by $\eta^{\mathcal{V}}_X(x)(i) = v_i(x)$ for all $x \in X$ and $i \in I$. Put $F^{\mathcal{V}}(X) = \langle \eta^{\mathcal{V}}_X(X)\rangle_P$, the subalgebra of P generated by the image of $\eta^{\mathcal{V}}_X$. Then by the construction, $F^{\mathcal{V}}(X) \in \mathcal{V}$. Moreover, $F^{\mathcal{V}}(X)$ with $\eta^{\mathcal{V}}_X \colon X \to |F^{\mathcal{V}}(X)|$ is free over X in $\mathcal{V}$. □

Free algebras over X in $\mathcal{V}$ are defined uniquely up to isomorphism. Moreover:

Corollary 3.5. *Let $\mathcal{V}$ be a variety of Σ-algebras and X be an S-sorted set. If Σ-algebra $F^{\mathcal{V}}(X) \in \mathcal{V}$ with function $\eta^{\mathcal{V}}_X \colon X \to |F^{\mathcal{V}}(X)|$ is free over X in $\mathcal{V}$ then $F^{\mathcal{V}}(X) \in \mathcal{V}$ is generated by $\eta^{\mathcal{V}}_X(X)$.* □

Consequently, the free algebra is a quotient of the term algebra $T_\Sigma(X)$. Here is the exact characterisation of the congruence on $T_\Sigma(X)$ that yields the free algebra:

Corollary 3.6. *Given any variety $\mathcal{V}$ of Σ-algebras and S-sorted set X, consider Σ-algebra $F^{\mathcal{V}}(X) \in \mathcal{V}$ with function $\eta^{\mathcal{V}}_X \colon X \to |F^{\mathcal{V}}(X)|$ free over X in $\mathcal{V}$. For any terms $t, t' \in |T_\Sigma(X)|_s$, $s \in S$, $t_{F^{\mathcal{V}}(X)}(\eta^{\mathcal{V}}_X) = t'_{F^{\mathcal{V}}(X)}(\eta^{\mathcal{V}}_X)$ iff $\mathcal{V} \models \forall X. t = t'$.* □

Lemma 3.4 carries over to everywhere non-empty algebras, albeit with an additional assumption (the same proof applies):

Lemma 3.7. *Let $\mathcal{V}$ be an everywhere non-empty variety of Σ-algebras and X be an S-sorted set such that $T_\Sigma(X)$ is everywhere non-empty. Then there exists a Σ-algebra $F^{\mathcal{V}}(X) \in \mathcal{V}$ with a function $\eta^{\mathcal{V}}_X \colon X \to |F^{\mathcal{V}}(X)|$ that is free over X in $\mathcal{V}$, that is, such that for any algebra $A \in \mathcal{V}$ and function $v \colon X \to |A|$, there exists a unique Σ-homomorphism $v^\sharp \colon F^{\mathcal{V}}(X) \to A$ with $\eta^{\mathcal{V}}_X;v^\sharp = v$.*

Moreover, such an algebra $F^{\mathcal{V}}(X) \in \mathcal{V}$ is defined uniquely up to isomorphism, $F^{\mathcal{V}}(X) \in \mathcal{V}$ is generated by $\eta^{\mathcal{V}}_X(X)$, and for any terms $t, t' \in |T_\Sigma(X)|_s$, $s \in S$, $t_{F^{\mathcal{V}}(X)}(\eta^{\mathcal{V}}_X) = t'_{F^{\mathcal{V}}(X)}(\eta^{\mathcal{V}}_X)$ iff $\mathcal{V} \models t = t'$. □

Lemmas 3.4 and 3.7 do not require the class $\mathcal{V}$ to be closed under homomorphic images: it is sufficient to assume $\mathcal{V}$ to be a *quasi-variety*, that is, closed under products and subalgebras.

The free algebras in a variety may be "adjusted" to freely accommodate any algebra:

Lemma 3.8. *Given any variety $\mathcal{V}$ of Σ-algebras and Σ-algebra A, there exists a Σ-algebra $F^{\mathcal{V}}(A) \in \mathcal{V}$ with a (surjective) Σ-homomorphism $\eta^{\mathcal{V}}_A \colon A \to F^{\mathcal{V}}(A)$ that is free over A in $\mathcal{V}$, that is, such that for any algebra $B \in \mathcal{V}$ and homomorphism $h \colon A \to B$, there is a unique Σ-homomorphism $h^{\sharp\sharp} \colon F^{\mathcal{V}}(A) \to B$ with $\eta^{\mathcal{V}}_A;h^{\sharp\sharp} = h$.*
Proof sketch. *Consider the Σ-algebra $F^{\mathcal{V}}(|A|) \in \mathcal{V}$ with a function $\eta^{\mathcal{V}}_{|A|} \colon |A| \to |F^{\mathcal{V}}(|A|)|$ that is free over $|A|$ in $\mathcal{V}$. Define an S-sorted relation $\equiv\ \subseteq |A| \times |A|$ as follows, for $s \in S$:*

$$\equiv_s = \{\langle t_{F^{\mathcal{V}}(|A|)}(\eta^{\mathcal{V}}_{|A|}), t'_{F^{\mathcal{V}}(|A|)}(\eta^{\mathcal{V}}_{|A|})\rangle \mid t, t' \in |T_\Sigma(|A|)|_s, t_A(id_{|A|}) = t'_A(id_{|A|})\}.$$

Then $\equiv$ is a congruence on $F^{\mathcal{V}}(|A|) \in \mathcal{V}$. Let $F^{\mathcal{V}}(A) = F^{\mathcal{V}}(|A|)/\equiv$ be the quotient algebra. Clearly, $F^{\mathcal{V}}(A) \in \mathcal{V}$. Consider the (surjective) function $\eta^{\mathcal{V}}_A \colon |A| \to |F^{\mathcal{V}}(A)|$ defined by $\eta^{\mathcal{V}}_A(a) = [\eta^{\mathcal{V}}_{|A|}(a)]_\equiv$. For any term $t \in |T_\Sigma(|A|)$, if $t_A(id_A) = a$ then $t_{F^{\mathcal{V}}(A)}(\eta^{\mathcal{V}}_A) = [t_{F^{\mathcal{V}}(|A|)}(\eta^{\mathcal{V}}_{|A|})]_\equiv = [\eta^{\mathcal{V}}_{|A|}(a)]_\equiv = \eta^{\mathcal{V}}_A(a)$. This proves that $\eta^{\mathcal{V}}_A$ is a Σ-homomorphism $\eta^{\mathcal{V}}_A \colon A \to F^{\mathcal{V}}(A)$. Now, given any homomorphism $h \colon A \to B$, for the homomorphism $h^\sharp \colon F^{\mathcal{V}}(|A|) \to B$ such that $\eta^{\mathcal{V}}_{|A|};h^\sharp = h$, we have $\equiv \subseteq ker(h^\sharp)$ (since for $t \in |T_\Sigma(|A|)|$, $h^\sharp(t_{F^{\mathcal{V}}(|A|)}(\eta^{\mathcal{V}}_{|A|})) = t_B(\eta^{\mathcal{V}}_{|A|};h^\sharp) = t_B(h) = h(t_A(id_{|A|}))$). Hence, we have a homomorphism $h^{\sharp\sharp} \colon F^{\mathcal{V}}(A) \to B$ defined by $h^{\sharp\sharp}([a]_\equiv) = h^\sharp(a)$, which is unique such that $\eta^{\mathcal{V}}_A;h^{\sharp\sharp} = h$. □

As for free algebras over a set, free algebras in a variety over an algebra are defined uniquely up to isomorphism.

Lemma 3.9. *Let $\mathcal{V}$ be a variety of Σ-algebras and A be a Σ-algebra, and let Σ-algebra $F^{\mathcal{V}}(A) \in \mathcal{V}$ with a homomorphism $\eta^{\mathcal{V}}_A \colon A \to F^{\mathcal{V}}(A)$ be free over A in $\mathcal{V}$. Then $A \in \mathcal{V}$ if and only if $\eta^{\mathcal{V}}_A$ is injective.*
Proof sketch. *Using Lemma 3.8: the "if" part follows since then $\eta^{\mathcal{V}}_A$ is bijective and A is isomorphic to $F^{\mathcal{V}}(A) \in \mathcal{V}$. For the "only if" part, just notice that if $A \in \mathcal{V}$ then $\eta^{\mathcal{V}}_A;id^{\sharp\sharp}_A$ is the identity.* □

The above lemma offers a criterion for checking whether or not an algebra A is in the variety considered. Given the explicit construction of $F^{\mathcal{V}}(A)$ and $\eta^{\mathcal{V}}_A \colon A \to F^{\mathcal{V}}(A)$ in Lemma 3.8, this has a crucial consequence:

Corollary 3.10. *Given any variety $\mathcal{V}$ of Σ-algebras and Σ-algebra A, if A is a model of $ETh(\mathcal{V})$ then $A \in \mathcal{V}$.*
Proof sketch. *By Lemma 3.9, it is enough to show that given the algebra $F^{\mathcal{V}}(A) \in \mathcal{V}$ with $\eta^{\mathcal{V}}_A \colon A \to F^{\mathcal{V}}(A)$ free over A in $\mathcal{V}$, $\eta^{\mathcal{V}}_A$ is injective. Consider any $a, a' \in |A|$ and suppose that $\eta^{\mathcal{V}}_A(a) = \eta^{\mathcal{V}}_A(a')$. Under the notation of the proof of Lemma 3.8, this means that $\eta^{\mathcal{V}}_{|A|}(a) \equiv \eta^{\mathcal{V}}_{|A|}(a')$. So, the definition of $\equiv$ implies that for some terms $t, t' \in |T_\Sigma(|A|)|$, we have $t_{F^{\mathcal{V}}_{|A|}}(\eta^{\mathcal{V}}_{|A|}) = \eta^{\mathcal{V}}_{|A|}(a)$, $t'_{F^{\mathcal{V}}_{|A|}}(\eta^{\mathcal{V}}_{|A|}) = \eta^{\mathcal{V}}_{|A|}(a')$, and $t_A(id_{|A|}) = t'_A(id_{|A|})$. Hence, by Cor. 3.6, equations $\forall |A|. t = a$ and $\forall |A|. t' = a'$ hold in $\mathcal{V}$, and so in A as well. Consequently, we have $a = t_A(id_{|A|}) = t'_A(id_{|A|}) = a'$, which completes the proof.* □

Finally, this yields the famous Birkhoff's variety theorem:

Theorem 3.11. *A class of Σ-algebras is equationally definable if and only if it is a variety.*
Proof sketch. *The "only if" part follows by Lemma 3.2. For the "if" part, consider any Σ-variety $\mathcal{V} \subseteq |$**Alg**$(\Sigma)|$. By definition we have $Mod(ETh(\mathcal{V})) \supseteq \mathcal{V}$. The opposite inclusion $Mod(ETh(\mathcal{V})) \subseteq \mathcal{V}$ follows by Cor. 3.10, which states that for any $A \in |$**Alg**$(\Sigma)|$ such that $A \models ETh(\mathcal{V})$, we also have $A \in \mathcal{V}$.* □

It is essential to consider arbitrary equations here, which in general may involve variables for an infinite number of sorts. Technically, the key step in the proof of Cor. 3.10 used equations which in general need not be finitary. A finitary variant of the same property requires an additional assumption:

Corollary 3.12. *Given any variety $\mathcal{V}$ of Σ-algebras and Σ-algebra A, if the set of sorts $\{s \in S \mid |A|_s \neq \emptyset\}$ is almost non-void in Σ and A is a model of $ETh_{fin}(\mathcal{V})$ then $A \in \mathcal{V}$.*

Proof sketch. *The proof of Cor. 3.10 carries over with additional remark that by Fact 2.1 and the assumptions, each of the equations $\forall |A|.\, t = a$ and $\forall |A|.\, t' = a'$ is equivalent to a finitary one.* □

This yields a finitary version of Birkhoff's variety theorem for some infinitary signatures:

Theorem 3.13. *Let Σ be such that all sets of sorts in Σ are almost non-void. Then a class of Σ-algebras is definable by finitary equations if and only if it is a variety.* □

Corollary 3.14. *Let Σ has a finite set of sort names. Then a class of Σ-algebras is definable by finitary equations if and only if it is a variety.* □

Unfortunately, the extra assumption about the signature cannot be dropped in general; here is a simple counter-example:

Example 3.15. *Consider a signature with an infinite set of sorts S and no operations. Let Φ be the set of equations of the form $\forall X \cup \{x, y\colon s\}.\, x = y$, where $s \in S$ and $X_{s'} \neq \emptyset$ for infinitely many sorts $s' \in S$. Then $Mod(\Phi)$ is a variety that consists of all Σ-algebras A such that either A is a subalgebra of $\mathbf{1}_\Sigma$ (i.e., for all sorts $s \in S$, $|A|_s$ has at most one element) or for almost all (i.e., all but finitely many) sorts $s \in S$, $|A|_s = \emptyset$. $Mod(\Phi)$ is not definable by finitary equations (for instance because it is not closed under directed sums).* □

In the example above, the set of all sorts is not almost non-void. However, in general this is a weaker property than required in Thm. 3.13: there are signatures where the set of all sorts is almost non-void, but some of its subsets are not. The example above can be adapted to show this requirement cannot be weakened.

Proposition 3.16. *Every Σ-variety is definable by finitary equations if and only if every set of sorts in Σ is almost non-void.*

Proof sketch. *Thm. 3.13 yields the "if" part. For the "only if" part, suppose that there is a set $S' \subseteq S$ of sorts that is not almost non-void. Let Φ be the set of equations of the form $\forall X \cup \{x, y\colon s\}.\, x = y$, where $s \in S$ and the set $\{s \in S \mid X_{s'} \neq \emptyset\}$ is not almost non-void. Then $Mod(\Phi)$ is a variety that consists of all Σ-algebras A such that either A is a subalgebra of $\mathbf{1}_\Sigma$ (i.e., for all sorts $s \in S$, $|A|_s$ has at most* *one element) or the set of sorts $\{s \in S \mid |A|_s \neq \emptyset\}$ is almost non-void. For instance, all algebras of the form $T_\Sigma(X_{S_0})$, where $S_0 \subseteq S'$ is a finite set of sorts and X_{S_0} contains exactly two variables (from some fixed vocabulary) of each sort $s \in S_0$ and is empty for all other sorts, satisfy Φ. But their direct sum has at least two elements for each sort in S', and so does not satisfy Φ. It follows that Mod(Φ) is not definable by finitary equations.* □

The last argument in Example 3.15 and the proof of Prop. 3.16 gives a crucial hint on the characterisation of varieties that are definable by finitary equations:

Theorem 3.17. *A class of Σ-algebras is definable by a set of finitary equations if and only if it is a finitary variety.*
Proof sketch. *The "only if" part follows by Lemma 3.2. For the "if" part, consider any Σ-variety $\mathcal{V} \subseteq |\mathbf{Alg}(\Sigma)|$ that is closed under directed sums. It is enough to show $Mod(ETh_{fin}(\mathcal{V})) \subseteq \mathcal{V}$. Consider $A \in Mod(ETh_{fin}(\mathcal{V}))$. Let $S_0 \subseteq S$ be a finite set of sorts, and let $|A|_{S_0}$ be the S-sorted set such that for all $s \in S$, $(|A|_{S_0})_s = |A|_s$ if $s \in S_0$ and $(|A|_{S_0})_s = \emptyset$ otherwise. Consider $A_{S_0} = \langle |A|_{S_0} \rangle_A$, the subalgebra of A generated by $|A|_{S_0}$. Since A is a model of $ETh_{fin}(\mathcal{V})$, so is A_{S_0} and hence by Cor. 3.12, $A_{S_0} \in \mathcal{V}$. Moreover, the family $\langle A_{S_0} \rangle_{S_0 \subseteq_{fin} S}$ (indexed by all finite $S_0 \subseteq_{fin} S$) is directed and A is its (directed) sum, $A = \coprod_{S_0 \subseteq_{fin} S} A_{S_0}$. Since $\mathcal{V}$ is closed under directed sums, we have $A \in \mathcal{V}$, which completes the proof.* □

None of the complications above arises when we deal with everywhere non-empty algebras. Given Lemma 3.7, we get variants of Lemmas 3.8 and 3.9, and Cor. 3.10 where $\mathcal{V}$ is an everywhere non-empty variety and A is an everywhere non-empty algebra. This yields:

Theorem 3.18. *A class of everywhere non-empty Σ-algebras is everywhere non-empty equationally definable (by a set of naive equations) if and only if it is an everywhere non-empty variety.* □

4 Equational calculus

Again, let $\Sigma = \langle S, \Omega \rangle$ be an algebraic signature.

The standard "naive" equational calculus is given in Table 1 (with obvious assumptions about the terms and operation names involved in the congruence rule, and the notation for substitution of terms for variables in terms introduced in Sect. 2.3). Given a set Φ of naive Σ-equations and a naive Σ-equation φ, we write $\Phi \vdash_{naive} \varphi$ whenever φ may be derived from Φ using the rules in Table 1.

It is easy to check that the naive equational calculus is sound for everywhere non-empty algebras:

Fact 4.1. *For any set Φ of naive Σ-equations and naive Σ-equation φ, if $\Phi \vdash^{naive} \varphi$ then $\Phi \models^{NE} \varphi$.* □

Reflexivity:	$\dfrac{}{t = t}$
Symmetry:	$\dfrac{t = t'}{t' = t}$
Transitivity:	$\dfrac{t = t' \quad t' = t''}{t = t''}$
Congruence:	$\dfrac{t_1 = t_1' \quad \cdots \quad t_n = t_n'}{f(t_1, \ldots, t_n) = f(t_1', \ldots, t_n')}$
Instantiation:	$\dfrac{t = t'}{t[\theta] = t'[\theta]}$

Table 1: Naive equational calculus

However, it may come as a surprise that the naive equational calculus is not sound for arbitrary many-sorted algebras:

Example 4.2. *Let signature Σ_0 have two sorts s, s', two constants $a, b\colon s$ and unary operation $f\colon s' \to s$. Then $\{f(x) = a, f(x) = b\} \vdash^{naive} a = b$. However, $\{f(x) = a, f(x) = b\} \not\models a = b$, since a Σ-algebra A with empty carrier $|A|_s = \emptyset$ and $a_A \neq b_A$ satisfies both $f(x) = a$ and $f(x) = b$, but does not satisfy $a = b$.*

The problem here is, of course, that $\{f(x) = a, f(x) = b\}$ semantically entail $\forall\{x\colon s'\}.\, a = b$, but when algebras with empty carriers are admitted, this does not entail $a = b$ (over the signature Σ_0). □

The example above gives a crucial hint on the essence of the problem, and indeed a sufficient and necessary condition for the soundness of the naive equational calculus may be easily formulated:

Proposition 4.3. *The naive equational calculus is sound if and only if for each operation name $f\colon s_1 \times \cdots \times s_n \to s$ in Σ, $\{s\}$ makes all sorts $s_1, \ldots, s_n$ non-void in Σ.*

Proof sketch. *For the "if" part, first check that all the rules in Table 1 except for transitivity are sound without additional assumptions on the signature. So, consider an instance of transitivity for some Σ-terms t, t', t'' of a common sort s. Using the assumption on the signature, we easily get that $\{s\}$ makes all the sorts of the variables in $FV(t')$ non-void in Σ. Now, given any Σ-algebra A such that $A \models t = t'$ and $A \models t' = t''$, any valuation $v\colon (FV(t) \cup FV(t'')) \to |A|$ extends to a valuation $v'\colon (FV(t) \cup FV(t'') \cup FV(t')) \to |A|$, and we have $t_A(v) = t_A(v') =$*

$t'_A(v') = t''_A(v') = t''_A(v)$, which proves that the "naive" rule of transitivity is sound under our assumption on Σ.

For the "only if" part, suppose that we have an operation $f\colon s_1 \times \cdots \times s_n \to s$ where for some $i = 1, \ldots, n$, $\{s\}$ does not make s_i non-void. Consider the term algebra $T_\Sigma(X)$, where X contains exactly two elements of sort s, and is empty for all other sorts. Then $|T_\Sigma(X)|_{s_i} = \emptyset$ and so (naive) Σ-equations $x = f(x_1, \ldots, x_n)$ and $f(x_1, \ldots, x_n) = y$ hold in $T_\Sigma(X)$, while $T_\Sigma(X) \not\models x = y$ (for any distinct x and y, and variables $x_1, \ldots, x_n$ of the appropriate sorts). □

In particular, this implies soundness of the naive equational calculus in the single-sorted context, even if we admit algebras with empty carriers:

Corollary 4.4. *Let Σ be a single-sorted signature. For any set Φ of naive Σ-equations and naive Σ-equation φ, if $\Phi \vdash^{naive} \varphi$ then $\Phi \models \varphi$.* □

The trick to make the equational calculus sound in the many-sorted context is to carefully keep track of the set of variables in the equations considered.

The resulting many-sorted equational calculus is given in Table 2 (with obvious assumptions about the terms and operation names involved in the congruence rule, with substitution $\theta\colon X \to |T_\Sigma(Y)|$ and the notation for substitution of terms for variables in terms in the instantiation rule). Given a set Φ of Σ-equations and a Σ-equation φ, we write $\Phi \vdash \varphi$ whenever φ may be derived from Φ using the rules in Table 2.

Reflexivity:	$\dfrac{}{\forall X.\, t = t}$
Symmetry:	$\dfrac{\forall X.\, t = t'}{\forall X.\, t' = t}$
Transitivity:	$\dfrac{\forall X.\, t = t' \quad \forall X.\, t' = t''}{\forall X.\, t = t''}$
Congruence:	$\dfrac{\forall X.\, t_1 = t'_1 \quad \cdots \quad \forall X.\, t_n = t'_n}{\forall X.\, f(t_1,' \ldots, t_n) = f(t'_1, \ldots, t'_n)}$
Instantiation:	$\dfrac{\forall X.\, t = t'}{\forall Y.\, t[\theta] = t'[\theta]}$

Table 2: Many-sorted equational calculus

The set of variables in the equations considered may be manipulated only by the instantiation rule. Apart from the obvious role (to substitute terms for variables that actually occur in the equated terms) this rule can also be used to remove variables that do not occur in these terms. This is possible when we have a term to be "substituted" for each such a variable, that is, if the sorts of the variables to be removed are made non-void by the set of sorts of the variables that remain in the equation. It can be easily checked that removing such variables is sound, even if many-sorted algebras with some carriers empty are considered. Indeed, the many-sorted calculus of Table 2 is sound:

Fact 4.5. *For any set Φ of Σ-equations and Σ-equation φ, if $\Phi \vdash \varphi$ then $\Phi \models \varphi$.* □

The many-sorted equational calculus as given in Table 2 is also complete:

Theorem 4.6. *Given a set Φ of Σ-equations and a Σ-equation φ, $\Phi \vdash \varphi$ if and only if $\Phi \models \varphi$.*

Proof sketch. *The "only if" part is Fact 4.5. For the "if" part, consider any set Φ of Σ-equations and Σ-equation $\forall X. t = t'$ such that $\Phi \models \forall X. t = t'$. Consider also the term algebra $T_\Sigma(X)$ and a binary relation $\approx$ on its carrier such that $t_1 \approx t_2$ iff $\Phi \vdash \forall X. t_1 = t_2$. It is easy to check that $\approx$ is a congruence on $T_\Sigma(X)$, so we can consider the quotient algebra $T_{\Sigma,\Phi}(X) = T_\Sigma(X)/\approx$.*

Now, for any S-sorted set Y, term $t_1 \in |T_\Sigma(Y)|$ and valuation $v\colon Y \to |T_{\Sigma,\Phi}(X)|$ with substitution $\theta_v\colon Y \to |T_\Sigma(X)|$ such that $v(y) = [\theta(y)]_\approx$ for all $y \in Y$, we have $(t_1)_{T_{\Sigma,\Phi}(X)}(v) = [t_1[\theta_v]]_\approx$. Hence, for any equation $\forall Y. t_1 = t'_1$ in Φ and valuation $v\colon Y \to |T_{\Sigma,\Phi}(X)|$ with substitution $\theta_v\colon Y \to |T_\Sigma(X)|$ as above, $(t_1)_{T_{\Sigma,\Phi}(X)}(v) = [t_1[\theta_v]]_\approx = [t'_1[\theta_v]]_\approx = (t_1)_{T_{\Sigma,\Phi}(X)}(v)$ since $\Phi \vdash \forall X. t_1[\theta_v] = t'_1[\theta_v]$ by the rule of instantiation. It follows that $T_{\Sigma,\Phi}(X)$ satisfies all equations in Φ, and so $T_{\Sigma,\Phi}(X) \models \forall X. t = t'$. Consequently, $[t]_\approx = t_{T_{\Sigma,\Phi}}(id_X) = t'_{T_{\Sigma,\Phi}}(id_X) = [t']_\approx$, and so $t \approx t'$. By definition this means that $\Phi \vdash \forall X. t = t'$, which completes the proof. □

An easy comparison of the calculi in Tables 1 and 2, respectively, shows that for any set Φ of naive equations, if $\Phi \vdash \forall X. t = t'$ then also $\Phi \vdash^{naive} t = t'$. Thus, Theorem 4.6 (and Fact 4.1) implies completeness of the naive calculus for everywhere non-empty semantic consequence:

Corollary 4.7. *Given a set Φ of naive Σ-equations and a naive Σ-equation φ, $\Phi \vdash^{naive} \varphi$ if and only if $\Phi \models^{NE} \varphi$.* □

By the same argument, but using Prop. 4.3, we get:

Corollary 4.8. *If (and only if) for each operation name $f\colon s_1 \times \cdots \times s_n \to s$ in Σ, $\{s\}$ makes all sorts $s_1, \ldots, s_n$ non-void in Σ then for any set Φ of naive Σ-equations and naive Σ-equation φ, $\Phi \vdash^{naive} \varphi$ if and only if $\Phi \models \varphi$.* □

5 Interpolation

The well-known interpolation property for first-order logic may be stated as follows [4]: for any sentences φ over signature Σ_1 and ψ over signature Σ_2 such that $\varphi \models \psi$ (over signature $\Sigma_1 \cup \Sigma_2$) there is a sentence δ (the interpolant) over the signature $\Sigma_1 \cap \Sigma_2$ such that $\varphi \models \delta$ (over signature Σ_1) and $\delta \models \psi$ (over signature Σ_2). In a sense, this is the interpolation property over the following (pushout) diagram of inclusions:

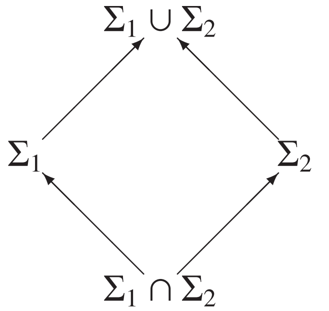

It is only natural to generalise this property to an arbitrary pushout square, as put forward in [16]. What is surprising then is that interpolation property for first-order logic fails in general, and holds only under additional assumptions about the signature morphisms involved.

Theorem 5.1. *Given a pushout in the category of signatures:*

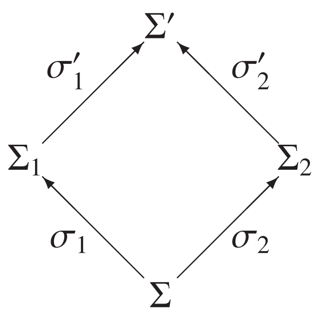

such that at least one of σ_1 or σ_2 is injective on sorts, for any first-order Σ_1-sentence φ and first-order Σ_2-sentence ψ such that $\sigma'_1(\varphi) \models_{\Sigma'} \sigma'_2(\psi)$, there exists a first-order Σ-sentence δ such that $\varphi \models_{\Sigma_1} \sigma_1(\delta)$ and $\sigma_2(\delta) \models_{\Sigma_2} \psi$. □

The proof of this property is beyond the scope of this paper, but see [3] (also for a counterexample showing that if both morphisms are non-injective on sorts then the interpolation property as formulated above may fail).

The equational logic does not have the interpolation property in the above form; however, it is often claimed that it does have an interpolation property in a version where sets of sentences and sets of interpolants are allowed. Unfortunately, this is not true in general in the many-sorted context, even if we limit attention to the pushouts of signature inclusions, as in the first formulation above:

Example 5.2. *Let Σ be a signature with three sorts s, s_1 and s_2, and two constants $a, b\colon s$. Let Σ_1 and Σ_2 extend Σ by a constant $e\colon s_1$ and by a unary operation*

$f\colon s_1 \to s_2$ respectively. Then $\Sigma = \Sigma_1 \cap \Sigma_2$ and let $\Sigma' = \Sigma_1 \cup \Sigma_2$. Consider Σ_1-equation $\forall\{x\colon s_2\}.\, a = b$ and Σ_2-equation $a = b$. Clearly, $\forall\{x\colon s_2\}.\, a = b \models_{\Sigma'} a = b$ (since all Σ'-algebras are everywhere non-empty).

Suppose that we have a set of Σ-equations Δ such that $\forall\{x\colon s_2\}.\, a = b \models_{\Sigma_1} \Delta$ and $\Delta \models_{\Sigma_2} a = b$. Consider a Σ_1-algebra A_1 with the carrier of sort s_2 empty and with $a_{A_1} \neq b_{A_1}$. Clearly then, $A_1 \models_{\Sigma_1} \forall\{x\colon s_2\}.\, a = b$, so $A_1 \models_{\Sigma_1} \Delta$, and consequently $A_1|_\Sigma \models_\Sigma \Delta$ as well. Take a subalgebra of $A_1|_\Sigma$ with the empty carrier of sort s_1, which satisfies Δ by Lemma 3.2, and consider its expansion A_2 to a Σ_2-algebra. Then $A_2 \models_{\Sigma_2} \Delta$ but $A_2 \not\models_{\Sigma_2} a = b$ — contradiction. □

The key problem here is caused again by algebras with some carriers empty. When everywhere non-empty algebras are considered, interpolation holds [14]:

Theorem 5.3. *Consider any algebraic signatures Σ_1 and Σ_2, set Φ of Σ_1-equations and set Ψ of Σ_2-equations such that $\Phi \models^{NE}_{\Sigma_1\cup\Sigma_2} \Psi$. Then there is a set Δ of $(\Sigma_1 \cap \Sigma_2)$-equations such that $\Phi \models^{NE}_{\Sigma_1} \Delta$ and $\Delta \models^{NE}_{\Sigma_2} \Psi$.*

Proof sketch. *Let $\Sigma = \Sigma_1 \cap \Sigma_2$ and $\Sigma' = \Sigma_1 \cup \Sigma_2$. Put $\Delta = ETh(Mod_{NE}(\Phi)|_\Sigma)$. Let $\mathcal{V} = Mod_{NE}(\Delta)$ be the least everywhere non-empty Σ-variety that contains $Mod_{NE}(\Phi)|_\Sigma$. Clearly, $\Phi \models^{NE}_{\Sigma_1} \Delta$.*

To show $\Delta \models^{NE}_{\Sigma_2} \Psi$, consider an everywhere non-empty Σ_2-algebra B_2 such that $B_2 \models_{\Sigma_2} \Delta$; we have then $B_2|_\Sigma \models_\Sigma \Delta$, and so $B_2|_\Sigma \in \mathcal{V}$. Since $Mod_{NE}(\Phi)$ is closed under products, and products are preserved by reduct functors, it follows by Fact 3.1 that $\mathcal{V}$ consists of all Σ-algebras B such that B is a homomorphic image of a subalgebra of A such that $A \in Mod_{NE}(\Phi)|_\Sigma$. Therefore, $B_2|_\Sigma$ is a homomorphic image of a subalgebra of $A_1|_\Sigma$ for some Σ_1-algebra $A_1 \in Mod_{NE}(\Phi)$. By Fact 2.2, there is an everywhere non-empty Σ_2-algebra A_2 such that $A_2|_\Sigma = A_1|_\Sigma$ and B_2 is a homomorphic image of a subalgebra of A_2. Then, by the amalgamation property, we have an everywhere non-empty Σ'-algebra A' such that $A'|_{\Sigma_1} = A_1$ and $A'|_{\Sigma_2} = A_2$. Since $A_1 \models_{\Sigma_1} \Phi$, it follows that $A' \models_{\Sigma'} \Phi$, so also $A' \models_{\Sigma'} \Psi$, and $A_2 \models_{\Sigma_2} \Psi$. Consequently, $B_2 \models_{\Sigma_2} \Psi$. □

The proof of Thm. 5.3 directly carries over to a somewhat more general formulation for a pushout of signature morphisms satisfying the condition which allow us to use the crucial Fact 2.2, see [13]:

Corollary 5.4. *Consider a pushout in the category of signatures:*

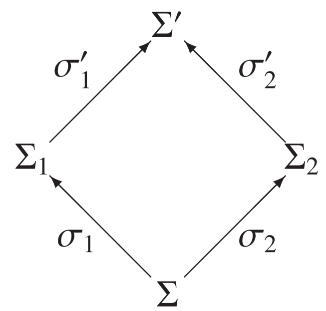

such that σ_2 is injective. For any sets Φ of Σ_1-equations and Ψ of Σ_2-equations such that $\sigma'_1(\Phi) \models^{NE}_{\Sigma'} \sigma'_2(\Psi)$, there is a set Δ of Σ-equations such that $\Phi \models^{NE}_{\Sigma_1} \sigma_1(\Delta)$ and $\sigma_2(\Delta) \models^{NE}_{\Sigma_2} \Psi$. □

The requirement in Cor. 5.4 that the "target" morphism σ_2 is injective cannot be dropped, not even for single-sorted signatures with non-void sort:

Example 5.5. *Let Σ be a signature with one sort s and constants $a, b, c_1, c_2 \colon s$, let Σ_1 be the extension of Σ by $f \colon s \to s$ with $\sigma_1 \colon \Sigma \to \Sigma_1$ being the inclusion, let Σ_2 be a signature with sort s and constants $a, b, c \colon s$, and let $\sigma_2 \colon \Sigma \to \Sigma_2$ map c_1 and c_2 to c and be identity otherwise. Consider the pushout of σ_1 and σ_2 as in Cor. 5.4.*

For Σ_1-equations $f(c_1) = a$ and $f(c_2) = b$ and Σ_2-equation $a = b$, we have $\sigma'_1(\{f(c_1) = a, f(c_2) = b\}) \models_{\Sigma'} \sigma'_2(a = b)$, since $\sigma'_1(c_1) = \sigma'_1(c_2)$. However, the least Σ-variety generated by $Mod(\{f(c_1) = a, f(c_2) = b\})\big|_{\sigma_1}$ contains all Σ-algebras ($\{f(c_1) = a, f(c_2) = b\}$ has only trivial equational consequences over Σ). Thus, there is no set Δ of Σ-equations such that $\{f(c_1) = a, f(c_2) = b\} \models_{\Sigma_1} \sigma_1(\Delta)$ and $\sigma_2(\Delta) \models_{\Sigma_2} a = b$. □

6 Final Remarks

We have presented some basic concepts and results of universal algebra in the many-sorted (heterogeneous) versions, recalling the standard Birkhoff's variety theorem, equational calculus and interpolation results, and their formulations in the many-sorted framework. We do not claim any technical originality here: the results are standard and known either from the literature or in the folklore (with some perhaps sharper then usual formulations following by inspection of the well-known proofs). So, instead of trying to summarise them here again, let us stress that *in essence* all the standard results of single-sorted (homogeneous) universal algebra seem to carry over to the many-sorted framework. However, their exact proper formulations require some care and adjustment:

- Birkhoff's variety theorem does not hold in general in the many-sorted framework when we allow signatures with infinite set of sorts, somewhere empty algebras (i.e., algebras with empty carriers of some sorts) and consider equations with finite set of variables only.
- Naive equational calculus is not sound for algebras that may be somewhere empty, unless some additional restrictions on signatures are considered.
- Equational interpolation property does not hold if somewhere empty algebras are considered.

One may wonder if we really want to consider signatures with infinite set of sorts, algebras with empty carriers, or signatures where some sorts are void. The point is that the need for such "anomalies" arises when we model some real phenomena in computer science. For instance, infinite signatures are needed when we want to consider polymorphic type systems, or any other type systems where the set of types is infinite (even if it may be finitely presented). Signatures with some sorts void are needed when we model genericity and parametrisation, either of specifications, or of programming modules. Finally, assuming all algebras to be everywhere non-empty not only makes a number of technical results more cumbersome (we loose the existence of reachable and initial algebras then, for instance) but again, excludes some natural models of real systems (for instance, consider database systems initialised with the empty set of data of some sort).

Consequently, instead of adopting any *ad hoc* assumptions, however standard and unproblematic they seem at first, the good practice is to formulate the concepts and results, even those that simply restate the standard ones, with sufficient care to take into account various nuances the many-sorted framework may bring.

References

[1] Garrett Birkhoff. On the structure of abstract algebras. *Proceedings of the Cambridge Philosophical Society*, 31:433–454, 1935.

[2] Garrett Birkhoff and John Lipson. Heterogeneous algebras. *Journal of Combinatorial Theory*, 8(1):115–133, 1970.

[3] Tomasz Borzyszkowski. Generalized interpolation in first order logic. *Fundamenta Informaticae*, 66(3):199–219, 2005.

[4] Chen-Chung Chang and H. Jerome Keisler. *Model Theory*. North-Holland, third edition, 1990.

[5] Hartmut Ehrig, Jacques Loeckx, and Bernd Mahr. A remark on the equational calculus for many-sorted algebras with possibly empty carrier sets. *Bulletin of the EATCS*, 30:61–65, 1986.

[6] Hartmut Ehrig and Bernd Mahr. *Fundamentals of Algebraic Specification 1*, volume 6 of *EATCS Monographs on Theoretical Computer Science*. Springer, 1985.

[7] Joseph Goguen and José Meseguer. Completeness of many sorted equational deduction. *Houston Journal of Mathematics*, 11(3):307–334, 1985.

[8] Joseph A. Goguen and José Meseguer. Remarks on remarks on many-sorted algebras with possibly emtpay carrier sets. *Bulletin of the EATCS*, 30:66–73, 1986.

[9] George A. Grätzer. *Universal Algebra*. Springer, second edition, 1979.

[10] Joseph Goguen, James Thatcher, and Eric Wagner. An initial algebra approach to the specification, correctness and implementation of abstract data types. Technical Report RC 6487, IBM Watson Research Center, Yorktown Heights NY, 1976. Also in: *Current Trends in Programming Methodology. Volume IV (Data Structuring)* (ed. R.T. Yeh), Prentice-Hall, 80–149, 1978.

[11] Phillip J. Higgins. Algebras with a scheme of operators. *Mathematische Nachrichten*, 27:115–132, 1963.

[12] José Meseguer and Joseph Goguen. Initiality, induction and computability. In Maurice Nivat and John C. Reynolds, editors, *Algebraic Methods in Semantics*, pages 459–541. Cambridge, 1985.

[13] Grigore Rosu and Joseph A. Goguen. On equational Craig interpolation. *Journal of Universal Computer Science*, 6(1):194–200, 2000.

[14] Pieter Hendrik Rodenburg. A simple algebraic proof of the equational interpolation theorem. *Algebra Universalis*, 28:48–51, 1991.

[15] Donald Sannella and Andrzej Tarlecki. *Foundations of Algebraic Specification and Formal Software Development*. Monographs in Theoretical Computer Science. An EATCS Series. Springer, 2011, to appear.

[16] Andrzej Tarlecki. Bits and pieces of the theory of institutions. In David H. Pitt, Samson Abramsky, Axel Poigné, and David E. Rydeheard, editors, *Proceedings of the Tutorial and Workshop on Category Theory and Computer Programming*, volume 240 of *Lecture Notes in Computer Science*, pages 334–360. Springer, 1986.

[17] Martin Wirsing. Algebraic specification. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science. Volume B (Formal Models and Semantics)*, pages 675–788. North-Holland and MIT Press, 1990.