

Chapter 2

Simple equational specifications

A specification is an unambiguous description of a signature Σ and a class of Σ -algebras. Because we model programs as algebras, a specification amounts to a characterisation of a class of programs. Each of these programs is regarded as a correct realisation of the specification.

Given a signature Σ (which, if finite, may be presented by simply listing its sort names and its operation names with their arities and result sorts), there are two basic techniques that may be used for describing a class of Σ -algebras. The first is to simply give a list of all the algebras in the class. Unfortunately, we are almost always interested in *infinite* classes of algebras, where this technique is useless (although sometimes this may be made to work if we can present a finite number of algebras that in some precise way represent the entire class we want). The second is to describe the functional behaviour of the algebras in the class by listing the properties (axioms) they are to satisfy. This is the fundamental specification technique used in work on algebraic specification and the one that will be studied in this chapter. The simplest and most common case is when the properties are expressed in the form of universally quantified equations; in most of this chapter, we restrict attention to this case. Section 2.7 indicates other forms of axioms that may be of use, along with some possible variations on the definitions of Chapter 1, and further possibilities will be discussed in Chapter 4. Since most of the results in this chapter are fairly standard and proofs are readily available in the literature, most proofs are left as exercises for the reader.

Chapters 5 and 8 will cover additional techniques for describing classes of algebras. All of these involve taking a class of algebras and performing a simple operation to obtain another class of algebras, often over a different signature. Using such methods, complex specifications of classes of complex algebras may be built from small and easily understood units.

2.1 Equations

Any given signature characterises the class of algebras over that signature. Although this fixes the names of sorts and operations, it is an exceedingly limited form of description since each such class contains a wide range of different algebras. Any two algebras taken from such a class may have carrier sets of different cardinalities and containing different elements; even if both algebras happen to have “matching” carrier sets, the results produced by applying operations may differ. For most applications it is necessary to focus on a subclass of algebras, obtained by imposing *axioms* which serve as constraints on the permitted behaviour of operations. One particularly simple form of axioms is equations, which constrain behaviour by asserting that two given terms have the same value. Equations have limited expressive power, but this disadvantage is to some extent balanced by the simplicity and convenience of reasoning in equational logic (see Sections 2.4 and 2.6).

Variables in equations will be taken from a fixed but arbitrary infinite set \mathcal{X} . We require \mathcal{X} to be closed under finite disjoint union: if $\langle X_i \rangle_{i \in I}$ is finite and $X_i \subseteq \mathcal{X}$ for all $i \in I$, then $\bigsqcup \langle X_i \rangle_{i \in I} \subseteq \mathcal{X}$. We use variable names like x, y, z in examples, and so we assume that these are all in \mathcal{X} . Throughout this section, let $\Sigma = \langle S, \Omega \rangle$ be a signature.

Definition 2.1.1 (Equation). A Σ -equation $\forall X \bullet t = t'$ consists of:

- a finite S -sorted set X (of variables), such that $X_s \subseteq \mathcal{X}$ for all $s \in S$; and
- two Σ -terms $t, t' \in |T_\Sigma(X)|_s$ for some sort $s \in S$.

A Σ -equation of the form $\forall \emptyset \bullet t = t'$ is called a *ground (Σ -)equation*, and will sometimes be written $t = t'$. \square

The explicit quantification over X in a Σ -equation $\forall X \bullet t = t'$ is essential, as will become clear in Section 2.4. It is nevertheless common practice to leave quantification implicit, writing $t = t'$ in place of $\forall FV(t) \cup FV(t') \bullet t = t'$, but we will not follow this practice except for ground equations.

Definition 2.1.2 (Satisfaction). A Σ -algebra A *satisfies* (or *is a model of*) a Σ -equation $\forall X \bullet t = t'$, written $A \models_\Sigma \forall X \bullet t = t'$, if for every (S -sorted) function $v: X \rightarrow |A|$, $t_A(v) = t'_A(v)$.

A satisfies (or is a model of) a set \mathcal{E} of Σ -equations, written $A \models_\Sigma \mathcal{E}$, if $A \models_\Sigma e$ for every equation $e \in \mathcal{E}$. A class \mathcal{A} of Σ -algebras satisfies a Σ -equation e , written $\mathcal{A} \models_\Sigma e$, if $A \models_\Sigma e$ for every $A \in \mathcal{A}$. Finally, a class \mathcal{A} of Σ -algebras satisfies a set \mathcal{E} of Σ -equations, written $\mathcal{A} \models_\Sigma \mathcal{E}$, if $A \models_\Sigma \mathcal{E}$ for every $A \in \mathcal{A}$ (equivalently, if $\mathcal{A} \models_\Sigma e$ for every $e \in \mathcal{E}$, i.e. $A \models_\Sigma e$ for every $A \in \mathcal{A}$ and $e \in \mathcal{E}$). \square

Notation. We sometimes write \models in place of \models_Σ when Σ is obvious. \square

Occasionally we will say that an equation e *holds* in an algebra A when $A \models e$, and similarly for sets of equations and classes of algebras.

Exercise 2.1.3. Recall $\Sigma 1$ and $A 1$ from Example 1.2.4. Give some $\Sigma 1$ -equations (both ground and non-ground) that are satisfied by $A 1$. Give some $\Sigma 1$ -equations (both ground and non-ground) that are *not* satisfied by $A 1$. \square

Exercise 2.1.4. If $\forall X \bullet t = t'$ is a Σ -equation and $X \subseteq X'$ (and $X'_s \subseteq \mathcal{X}$ for all $s \in S$), it follows from Definition 2.1.1 that $\forall X' \bullet t = t'$ is also a Σ -equation. Show that $A \models_{\Sigma} \forall X \bullet t = t'$ implies that $A \models_{\Sigma} \forall X' \bullet t = t'$. Give a counterexample showing that the converse does *not* hold. (HINT: Consider $X_s = \emptyset$ and $|A|_s = \emptyset$ for some $s \in S$.) Show that it *does* hold if Σ has only one sort. \square

Exercise 2.1.5. Show that surjective Σ -homomorphisms preserve satisfaction of Σ -equations: if $h: A \rightarrow B$ is a surjective Σ -homomorphism then $A \models_{\Sigma} e$ implies $B \models_{\Sigma} e$ for any Σ -equation e . Show that injective Σ -homomorphisms reflect satisfaction of Σ -equations: if $h: A \rightarrow B$ is an injective Σ -homomorphism then $B \models_{\Sigma} e$ implies $A \models_{\Sigma} e$ for any Σ -equation e . Conclude that Σ -isomorphisms preserve and reflect satisfaction of Σ -equations. \square

Exercise 2.1.6. Give an alternative definition of $A \models_{\Sigma} \forall X \bullet t = t'$ via the satisfaction of $t = t'$ viewed as a ground equation over an enlarged signature. HINT: Definition 2.1.2 involves quantification over valuations $v: X \rightarrow |A|$. Consider how this might be replaced by quantification over algebras having a signature obtained from Σ by adding a constant for each variable in X . \square

A signature morphism $\sigma: \Sigma \rightarrow \Sigma'$ gives rise to a translation of Σ -equations to Σ' -equations. This is essentially a simple matter of applying the translation on terms induced by σ to both sides of the equation.

Definition 2.1.7 (Equation translation). Let $\forall X \bullet t = t'$ be a Σ -equation, and let $\sigma: \Sigma \rightarrow \Sigma'$ be a signature morphism. Recall from Definition 1.5.10 that we then have $\sigma(t), \sigma(t') \in |T_{\Sigma'}(X')|$ where

$$X'_{s'} = \bigsqcup_{\sigma(s)=s'} X_s \quad \text{for each } s' \in S'.$$

The translation of $\forall X \bullet t = t'$ by σ is then the Σ' -equation $\sigma(\forall X \bullet t = t') = \forall X' \bullet \sigma(t) = \sigma(t')$. (The fact that \mathcal{X} is closed under finite disjoint union guarantees that this is indeed a Σ' -equation.) \square

An important result which brings together some of the main definitions above is the following:

Lemma 2.1.8 (Satisfaction Lemma [BG80]). *If $\sigma: \Sigma \rightarrow \Sigma'$ is a signature morphism, e is a Σ -equation and A' is a Σ' -algebra, then $A' \models_{\Sigma'} \sigma(e)$ iff $A'|_{\sigma} \models_{\Sigma} e$.* \square

When e is a ground Σ -equation, it is easy to see that this follows directly from the property established in Exercise 1.5.12. When σ is injective (on both sort and operation names), it seems intuitively clear that the Satisfaction Lemma should hold, since the domain of quantification of variables is unchanged, the only difference between e and $\sigma(e)$ is the names used for sorts and operations, and the only difference between A' and $A'|_{\sigma}$ (apart from sort/operation names) is that A' might provide interpretations for sort and operation names which do not appear in $\sigma(e)$ and so cannot affect its satisfaction. When σ is non-injective the Satisfaction Lemma still holds, but this is less intuitively obvious (particularly when σ is non-injective on sort names).

Exercise 2.1.9. Take a signature morphism $\sigma: \Sigma \rightarrow \Sigma'$ which is non-injective on sort and operation names, a Σ -equation involving the sort and operation names for which σ is not injective, and a Σ' -algebra, and check that the Satisfaction Lemma holds in this case. \square

Exercise 2.1.10. Prove the Satisfaction Lemma, using Exercise 1.5.12. \square

Exercise 2.1.11. Define the translation of a Σ -equation by a derived signature morphism $\delta: \Sigma \rightarrow \Sigma'$, and convince yourself that the Satisfaction Lemma also holds for this case. \square

The Satisfaction Lemma says that the translations of syntax (terms, equations) and semantics (algebras) induced by signature morphisms are coherent with the definition of satisfaction. Said another way, the manner in which satisfaction of equations by algebras varies according to the signature at hand fits exactly with these translations. Further discussion of the property embodied in the Satisfaction Lemma may be found in Section 4.1.

2.2 Flat specifications

A signature together with a set of equations over that signature constitutes a simple form of specification. We refer to these as *flat* (meaning *unstructured*) specifications in order to distinguish them from the *structured* specifications to be introduced in Chapter 5, formed from simpler specifications using specification-building operations. As we shall see later, it is possible in some (but not all) cases to “flatten” a structured specification to yield a flat specification describing the same class of algebras. Throughout this section, let Σ be a signature.

Definition 2.2.1 (Presentation). A *presentation* (also known as a *flat specification*) is a pair $\langle \Sigma, \mathcal{E} \rangle$ where \mathcal{E} is a set of Σ -equations (called the *axioms* of $\langle \Sigma, \mathcal{E} \rangle$). A presentation $\langle \Sigma, \mathcal{E} \rangle$ is sometimes referred to as a Σ -*presentation*. \square

The term “presentation” is chosen to emphasize the syntactic nature of the concept. The idea is that a presentation *denotes* (or *presents*) a semantic object which is inconvenient to describe directly. A reasonable objection to the definition above is that it fails to include restrictions to ensure that presentations are truly syntactic objects, namely that Σ and \mathcal{E} are *finite*, or at least effectively presentable in some other sense (e.g. recursive or recursively enumerable). Although it would be possible to impose such a restriction, we refrain from doing so in order to avoid placing undue emphasis on issues of this kind.

Definition 2.2.2 (Model of a presentation). A *model* of a presentation $\langle \Sigma, \mathcal{E} \rangle$ is a Σ -algebra A such that $A \models_{\Sigma} \mathcal{E}$. $\text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$ is the class of all models of $\langle \Sigma, \mathcal{E} \rangle$. \square

Taking $\langle \Sigma, \mathcal{E} \rangle$ to denote the semantic object $\text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$ is sometimes called taking its *loose semantics*. The word “loose” here refers to the fact that this is not always

(in fact, hardly ever) an isomorphism class of algebras: $A, B \in \text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$ does *not* imply that $A \cong B$. In Section 2.5 we will consider the so-called *initial semantics* of presentations in which a further constraint is imposed on the models of a presentation, forcing every presentation to denote an isomorphism class of algebras.

Example 2.2.3. Let $\text{BOOL} = \langle \Sigma_{\text{BOOL}}, \mathcal{E}_{\text{BOOL}} \rangle$ be the presentation below.¹

spec $\text{BOOL} = \text{sorts } \text{Bool}$
ops $\text{true} : \text{Bool}$
 $\text{false} : \text{Bool}$
 $\neg _ : \text{Bool} \rightarrow \text{Bool}$
 $_ \wedge _ : \text{Bool} \times \text{Bool} \rightarrow \text{Bool}$
 $_ \Rightarrow _ : \text{Bool} \times \text{Bool} \rightarrow \text{Bool}$
 $\forall p, q : \text{Bool}$

- $\neg \text{true} = \text{false}$
- $\neg \text{false} = \text{true}$
- $p \wedge \text{true} = p$
- $p \wedge \text{false} = \text{false}$
- $p \wedge \neg p = \text{false}$
- $p \Rightarrow q = \neg(p \wedge \neg q)$

Define Σ_{BOOL} -algebras $A1$, $A2$ and $A3$ as follows:

$|A1|_{\text{Bool}} = \{\star\}$
 $\text{true}_{A1} = \star$
 $\text{false}_{A1} = \star$

$\neg_{A1} = \{\star \mapsto \star\}$

\wedge_{A1}	\star
\star	\star

\Rightarrow_{A1}	\star
\star	\star

$|A2|_{\text{Bool}} = \{\clubsuit, \heartsuit, \spadesuit\}$
 $\text{true}_{A2} = \clubsuit$
 $\text{false}_{A2} = \heartsuit$

$\neg_{A2} = \{\clubsuit \mapsto \heartsuit,$
 $\heartsuit \mapsto \clubsuit,$
 $\spadesuit \mapsto \spadesuit\}$

\wedge_{A2}	\clubsuit	\heartsuit	\spadesuit
\clubsuit	\clubsuit	\heartsuit	\heartsuit
\heartsuit	\heartsuit	\heartsuit	\heartsuit
\spadesuit	\spadesuit	\heartsuit	\heartsuit

\Rightarrow_{A2}	\clubsuit	\heartsuit	\spadesuit
\clubsuit	\clubsuit	\heartsuit	\clubsuit
\heartsuit	\clubsuit	\clubsuit	\clubsuit
\spadesuit	\clubsuit	\spadesuit	\clubsuit

$|A3|_{\text{Bool}} = \{tt, ff\}$
 $\text{true}_{A3} = tt$
 $\text{false}_{A3} = ff$

$\neg_{A3} = \{tt \mapsto ff,$
 $ff \mapsto tt\}$

\wedge_{A3}	tt	ff
tt	tt	ff
ff	ff	ff

\Rightarrow_{A3}	tt	ff
tt	tt	ff
ff	tt	tt

¹ Here and in the sequel we use notation from OBJ [KKM88] and CASL [Mos04] to introduce infix, prefix and “mixfix” operations. We also follow CASL by itemizing axioms in specifications, marking them with • and introducing universal quantification over the variables only once for the entire list of axioms. Although the meaning of an axiom can be affected by adding quantification over variables that it does not contain — see Exercise 2.1.4 — this pathology does not arise in any of our examples.

Each of these algebras is a model of **BOOL**. (NOTE: Reference will be made to **BOOL** and to its models *A1*, *A2* and *A3* in later sections of this chapter. The name **BOOL** has been chosen for the same reason as `bool` is used for the type of truth values in programming languages; it is technically a misnomer since this is not a specification of Boolean algebras; see Example 2.2.4 below.)

Exercise. Show that the models defined and in fact all the models of **BOOL** satisfy $\forall p:Bool \bullet \neg(p \wedge \neg false) = \neg p$. Define a model of **BOOL** that does not satisfy $\forall p:Bool \bullet \neg\neg p = p$. \square

Example 2.2.4. Let $BA = \langle \Sigma_{BA}, \mathcal{E}_{BA} \rangle$ be the following presentation.

```

spec BA = sorts Bool
           ops true: Bool
                false: Bool
                ¬_ : Bool → Bool
                _∨_ : Bool × Bool → Bool
                _∧_ : Bool × Bool → Bool
                _⇒_ : Bool × Bool → Bool
            $\forall p, q, r: Bool$ 
           •  $p \vee (q \wedge r) = (p \vee q) \wedge r$ 
           •  $p \wedge (q \vee r) = (p \wedge q) \vee r$ 
           •  $p \vee q = q \vee p$ 
           •  $p \wedge q = q \wedge p$ 
           •  $p \vee (p \wedge q) = p$ 
           •  $p \wedge (p \vee q) = p$ 
           •  $p \vee (q \wedge r) = (p \vee q) \wedge (p \vee r)$ 
           •  $p \wedge (q \vee r) = (p \wedge q) \vee (p \wedge r)$ 
           •  $p \vee \neg p = \text{true}$ 
           •  $p \wedge \neg p = \text{false}$ 
           •  $p \Rightarrow q = \neg p \vee q$ 

```

Models of **BA** are called *Boolean algebras*. One such model is the following two-valued Boolean algebra \mathbb{B} :

```

| $\mathbb{B}$ |Bool = {tt, ff},
true $\mathbb{B}$  = tt,
false $\mathbb{B}$  = ff,
¬ $\mathbb{B}$  = {tt ↦ ff, ff ↦ tt}

```

and

∨ \mathbb{B}	tt	ff
tt	tt	tt
ff	tt	ff

∧ \mathbb{B}	tt	ff
tt	tt	ff
ff	ff	ff

⇒ \mathbb{B}	tt	ff
tt	tt	ff
ff	tt	tt

This is essentially the same as *A3* in Example 2.2.3. Note that *A1* can be turned into a (trivial) Boolean algebra in a similar way, but this is not the case with *A2*.

Exercise. Given a Boolean algebra B , define a relation $\leq_B \subseteq |B| \times |B|$ by $a \leq_B b$ iff $a \vee_B b = b$. Show that \leq_B is a partial order with true_B and false_B as its greatest and least elements respectively, and with $a \vee_B b$ yielding the least upper bound of a, b and $a \wedge_B b$ yielding their greatest lower bound. (In fact, $\langle |B|, \leq_B \rangle$ is a distributive lattice with top and bottom elements and complement \neg_B .) \square

Exercise 2.2.5. Show that all Boolean algebras (the models of BA as introduced in Exercise 2.2.4) satisfy the *de Morgan laws*:

$$\begin{aligned} \forall p, q: \text{Bool} \bullet \neg(p \vee q) &= \neg p \wedge \neg q \\ \forall p, q: \text{Bool} \bullet \neg(p \wedge q) &= \neg p \vee \neg q. \end{aligned} \quad \square$$

The following characterisation of the expressive power of flat equational specifications is one of the classical theorems of universal algebra.

Definition 2.2.6 (Equationally definable class). A class \mathcal{A} of Σ -algebras is *equationally definable* if $\mathcal{A} = \text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$ for some set \mathcal{E} of Σ -equations. \square

Definition 2.2.7 (Variety). A class \mathcal{A} of Σ -algebras is *closed under subalgebras* if for any $A \in \mathcal{A}$ and subalgebra B of A , $B \in \mathcal{A}$. Similarly, \mathcal{A} is *closed under homomorphic images* if for any $A \in \mathcal{A}$ and Σ -homomorphism $h: A \rightarrow B$, $h(A) \in \mathcal{A}$, and \mathcal{A} is *closed under products* if for any family $\langle A_i \in \mathcal{A} \rangle_{i \in I}$, $\prod \langle A_i \rangle_{i \in I} \in \mathcal{A}$.

A non-empty class of Σ -algebras which is closed under subalgebras, homomorphic images, and products is called a *variety*. \square

Proposition 2.2.8. Any equationally definable class \mathcal{A} of Σ -algebras is a variety. \square

Exercise 2.2.9. Prove Proposition 2.2.8: show that for any presentation $\langle \Sigma, \mathcal{E} \rangle$, $\text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$ is closed under subalgebras, homomorphic images and products. For example, formalise the following argument to show closure under subalgebras: if $A \models_{\Sigma} e$ and B is a subalgebra of A then $B \models_{\Sigma} e$ since removing values from the carriers of an algebra does not affect the truth of universally quantified assertions about its behaviour. Closure under products and under homomorphic images are not much more difficult to prove. \square

Theorem 2.2.10 (Birkhoff's Variety Theorem [Bir35]). If Σ is a signature with a finite set of sort names then a class \mathcal{A} of Σ -algebras is a variety iff \mathcal{A} is equationally definable. \square

The “if” part of this theorem is (a special case of) Proposition 2.2.8. A complete proof of the “only if” part is beyond the scope of this book; the curious reader should consult [Wec92].

Example 2.2.11. Consider the signature

$$\begin{aligned} \Sigma = & \text{sorts } s \\ & \text{ops } 0: s \\ & \quad \dots \times \dots: s \times s \rightarrow s \end{aligned}$$

and the class \mathcal{A} of Σ -algebras satisfying the familiar cancellation law:

if $a \neq 0$ and $a \times b = a \times c$ then $b = c$.

The Σ -algebra A such that $|A|_s$ is the set of natural numbers and \times_A is ordinary multiplication is in \mathcal{A} . The Σ -algebra B such that $|B|_s = \{0, 1, 2, 3\}$ and \times_B is multiplication modulo 4 is not in \mathcal{A} . (**Exercise:** Why not?) Since B is a homomorphic image of A , this shows that \mathcal{A} is not a variety and hence is not equationally definable. \square

Exercise 2.2.12. Formulate a definition of what it means for a class of Σ -algebras to be closed under homomorphic coimages. Are varieties closed under homomorphic coimages? \square

Exercise 2.2.13. Formulate definitions of what it means for a class of Σ -algebras to be closed under quotients and under isomorphisms. Show that closure under both quotients and isomorphisms is equivalent to closure under homomorphic images. \square

The assumption in Theorem 2.2.10 that the set of sort names in Σ is finite cannot easily be omitted:

Exercise 2.2.14. A family \mathcal{B} of Σ -algebras is *directed* if any two algebras $B_1, B_2 \in \mathcal{B}$ are subalgebras of some $B \in \mathcal{B}$. Define the *union* $\bigcup \mathcal{B}$ of such a family to be the least Σ -algebra such that each $B \in \mathcal{B}$ is a subalgebra of $\bigcup \mathcal{B}$ (the carrier of $\bigcup \mathcal{B}$ is the union of the carriers of all algebras in \mathcal{B} , and the values of operations on arguments are inherited from the algebras in \mathcal{B} ; this is well defined since \mathcal{B} is directed). Prove that since we consider equations with finite sets of variables only, then for any presentation $\langle \Sigma, \mathcal{E} \rangle$, $\text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$ is *closed under directed unions*, that is, given any *directed* family of algebras $\mathcal{B} \subseteq \text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$, its union $\bigcup \mathcal{B}$ is also in $\text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$.

A generalisation of Theorem 2.2.10 that we hint at here without a proof is that for *any* signature Σ , a class of Σ -algebras is equationally definable iff it is a variety that is closed under directed unions. \square

Exercise 2.2.15. Consider a signature with an infinite set of sort names and no operations. Let \mathcal{A}_{fn} be the class of all algebras over this signature that have non-empty carriers for a finite set of sorts only, and let \mathcal{A} be the closure of \mathcal{A}_{fn} under products and subalgebras (this adds algebras where the carrier of each sort is either a singleton or empty). Check that \mathcal{A} is a variety. Prove, however, that \mathcal{A} is not definable by any set of equations. HINT: Use Exercise 2.2.14. \square

Exercise 2.2.16. Modify the definition of equation (Definition 2.1.1) so that infinite sets of variables are allowed; it is enough to consider sets of variables that are finite for each sort, but may be non-empty for infinitely many sorts. Extend the notion of satisfaction (Definition 2.1.2) to such generalised equations in the obvious way. Check that the class \mathcal{A} defined in Exercise 2.2.15 is definable by such equations.

HINT: Consider all equations of the form $\forall X \cup \{x, y: s\} \bullet x = y$, for all sorts s and sets X of variables such that $X_{s'} \neq \emptyset$ for infinitely many sorts s' .

Another generalisation of Theorem 2.2.10 that we want to hint at here is that for *any* signature Σ a class of Σ -algebras is definable by such generalised equations iff it is a variety. The proof of the “if” part is as easy as for ordinary equations (Proposition 2.2.8). The proof of the “only if” part is quite similar to that of the finitary case. \square

A final remark to clarify the nuances in the many-sorted versions of Theorem 2.2.10 is that the theorem holds for *any* signature (also with an infinite set of sort names) when we restrict attention to algebras with non-empty carriers of all sorts: all varieties of such algebras (with closure under subalgebras limited to subalgebras with non-empty carriers) are definable by equations with a finite set of variables.

2.3 Theories

Any equationally definable class of algebras has many different presentations; in practice the choice of presentation is determined by various factors, including the need for simplicity and understandability and the desire for elegance. On the other hand, such a class uniquely determines the largest set of equations that defines it, called its theory. Since this is an infinite set, it is not a useful way of presenting the class. However, it is a useful set to consider since it contains all axioms in all presentations of the class, together with all their consequences.

Throughout this section, let Σ be a signature.

Definition 2.3.1 ($Mod_{\Sigma}(\mathcal{E})$, $Th_{\Sigma}(\mathcal{A})$, $Cl_{\Sigma}(\mathcal{E})$, $Cl_{\Sigma}(\mathcal{A})$). Given any set \mathcal{E} of Σ -equations, $Mod_{\Sigma}(\mathcal{E})$ (the *models of \mathcal{E}*) denotes the class of all Σ -algebras satisfying all the Σ -equations in \mathcal{E} :

$$Mod_{\Sigma}(\mathcal{E}) = \{A \mid A \text{ is a } \Sigma\text{-algebra and } A \models_{\Sigma} \mathcal{E}\} \quad (= Mod[\langle \Sigma, \mathcal{E} \rangle]).$$

For any class \mathcal{A} of Σ -algebras, $Th_{\Sigma}(\mathcal{A})$ (the *theory of \mathcal{A}*) denotes the set of all Σ -equations satisfied by each Σ -algebra in \mathcal{A} :

$$Th_{\Sigma}(\mathcal{A}) = \{e \mid e \text{ is a } \Sigma\text{-equation and } \mathcal{A} \models_{\Sigma} e\}.$$

A set \mathcal{E} of Σ -equations is *closed* if $\mathcal{E} = Th_{\Sigma}(Mod_{\Sigma}(\mathcal{E}))$. The *closure* of a set \mathcal{E} of Σ -equations is the (closed) set $Cl_{\Sigma}(\mathcal{E}) = Th_{\Sigma}(Mod_{\Sigma}(\mathcal{E}))$. Analogously, a class \mathcal{A} of Σ -algebras is *closed* if $\mathcal{A} = Mod_{\Sigma}(Th_{\Sigma}(\mathcal{A}))$, and the *closure* of \mathcal{A} is $Cl_{\Sigma}(\mathcal{A}) = Mod_{\Sigma}(Th_{\Sigma}(\mathcal{A}))$. \square

Proposition 2.3.2. For any sets \mathcal{E} and \mathcal{E}' of Σ -equations and classes \mathcal{A}, \mathcal{B} of Σ -algebras:

1. If $\mathcal{E} \subseteq \mathcal{E}'$ then $Mod_{\Sigma}(\mathcal{E}) \supseteq Mod_{\Sigma}(\mathcal{E}')$.

2. If $\mathcal{B} \supseteq \mathcal{A}$ then $Th_\Sigma(\mathcal{B}) \subseteq Th_\Sigma(\mathcal{A})$.
3. $\mathcal{E} \subseteq Th_\Sigma(Mod_\Sigma(\mathcal{E}))$ and $Mod_\Sigma(Th_\Sigma(\mathcal{A})) \supseteq \mathcal{A}$.
4. $Mod_\Sigma(\mathcal{E}) = Mod_\Sigma(Th_\Sigma(Mod_\Sigma(\mathcal{E})))$ and $Th_\Sigma(\mathcal{A}) = Th_\Sigma(Mod_\Sigma(Th_\Sigma(\mathcal{A})))$.
5. $Cl_\Sigma(\mathcal{E})$ and $Cl_\Sigma(\mathcal{A})$ are closed.

Proof. Exercise. HINT: Properties 4 and 5 follow from properties 1–3. □

For any signature Σ , the functions Th_Σ and Mod_Σ constitute what is known in lattice theory as a Galois connection.

Definition 2.3.3 (Galois connection). A *Galois connection* is given by two partially ordered sets A and M (in Proposition 2.3.2, A is the set of all sets of Σ -equations, and M is the “set” of all classes of Σ -algebras, both ordered by inclusion) and maps $_*: A \rightarrow M$ and $_+: M \rightarrow A$ (here Mod_Σ and Th_Σ) satisfying properties corresponding to 2.3.2(1)–2.3.2(3). An element $a \in A$ (or $m \in M$) is called *closed* if $a = (a^*)^+$ (or $m = (m^+)^*$). □

Some useful properties — including ones corresponding to 2.3.2(4) and 2.3.2(5) — hold for any Galois connection.

Exercise 2.3.4. For any Galois connection and any $a, b \in A$ and $m \in M$, show that the following properties hold:

1. $a \leq_A m^+$ iff $a^* \geq_M m$.
2. If a and b are closed then $a \leq_A b$ iff $a^* \geq_M b^*$. (Show that the “if” part fails if a or b is not closed.)

Here, \leq_A and \leq_M are the orders on A and M respectively. □

Exercise 2.3.5. For any Galois connection such that A and M have binary least upper bounds (\sqcup_A, \sqcup_M) and greatest lower bounds (\sqcap_A, \sqcap_M), and for any $a, b \in A$, show that the following properties hold:

1. $(a \sqcup_A b)^* = a^* \sqcap_M b^*$.
2. $(a \sqcap_A b)^* \geq_M a^* \sqcup_M b^*$.

HINT: \sqcup_A satisfies the following properties for any $a, b, c \in A$:

- $a \leq_A a \sqcup_A b$ and $b \leq_A a \sqcup_A b$.
- If $a \leq_A c$ and $b \leq_A c$ then $a \sqcup_A b \leq_A c$.

And analogously for \sqcap_A, \sqcup_M and \sqcap_M .

State and prove analogues to 1 and 2 for any $m, n \in M$, and instantiate all these general properties for the Galois connection between sets of Σ -equations and classes of Σ -algebras. □

Definition 2.3.6 (Semantic consequence). A Σ -equation e is a *semantic consequence* of a set \mathcal{E} of Σ -equations, written $\mathcal{E} \models_\Sigma e$, if $e \in Cl_\Sigma(\mathcal{E})$ (equivalently, if $Mod_\Sigma(\mathcal{E}) \models_\Sigma e$). □

Notation. We write $\mathcal{E} \models e$ instead of $\mathcal{E} \models_\Sigma e$ when the signature Σ is obvious. □

The use of the double turnstile (\models) here is the same as its use in logic: $\mathcal{E} \models e$ if the equation e is satisfied in every algebra which satisfies all the equations in \mathcal{E} . Here, \mathcal{E} is a set of *assumptions* and e is a *conclusion* which *follows from* \mathcal{E} . We refer to this as *semantic* (or *model-theoretic*) consequence to distinguish it from a similar relation defined by means of “syntactic” inference rules in the next section.

Example 2.3.7. Recall Example 2.2.3. The exercise there shows the following:

$$\begin{aligned} \mathcal{E}_{\text{Bool}} &\models_{\Sigma_{\text{Bool}}} \forall p:\text{Bool} \bullet \neg(p \wedge \neg \text{false}) = \neg p \\ \mathcal{E}_{\text{Bool}} &\not\models_{\Sigma_{\text{Bool}}} \forall p:\text{Bool} \bullet \neg\neg p = p \end{aligned}$$

Then, referring to Example 2.2.4, Exercise 2.2.5 shows that the de Morgan laws are semantical consequences of the set of axioms \mathcal{E}_{BA} . \square

Exercise 2.3.8. Prove that semantic consequence is preserved by translation along signature morphisms: for any signature morphism $\sigma: \Sigma \rightarrow \Sigma'$, set \mathcal{E} of Σ -equations, and Σ -equation e ,

$$\text{if } \mathcal{E} \models_{\Sigma} e \text{ then } \sigma(\mathcal{E}) \models_{\Sigma'} \sigma(e).$$

Equivalently, $\sigma(Cl_{\Sigma}(\mathcal{E})) \subseteq Cl_{\Sigma'}(\sigma(\mathcal{E}))$. Show that neither the reverse implication nor the reverse inclusion hold in general. \square

Exercise 2.3.9. Let $\sigma: \Sigma \rightarrow \Sigma'$ be a signature morphism and let \mathcal{E}' be a closed set of Σ' -equations. Show that $\sigma^{-1}(\mathcal{E}')$ is a closed set of Σ -equations. \square

See Section 4.2 for some further results on semantic consequence and translation along signature morphisms, presented in a more general context.

Definition 2.3.10 (Theory). A *theory* is a presentation $\langle \Sigma, \mathcal{E} \rangle$ such that \mathcal{E} is closed. A presentation $\langle \Sigma, \mathcal{E} \rangle$ (where \mathcal{E} need not be closed) *presents* the theory $\langle \Sigma, Cl_{\Sigma}(\mathcal{E}) \rangle$. A theory $\langle \Sigma, \mathcal{E} \rangle$ is sometimes referred to as a Σ -*theory*. \square

A theory morphism between two theories is a signature morphism between their signatures that maps the equations in the source theory to equations belonging to the target theory.

Definition 2.3.11 (Theory morphism). For any theories $\langle \Sigma, \mathcal{E} \rangle$ and $\langle \Sigma', \mathcal{E}' \rangle$, a *theory morphism* $\sigma: \langle \Sigma, \mathcal{E} \rangle \rightarrow \langle \Sigma', \mathcal{E}' \rangle$ is a signature morphism $\sigma: \Sigma \rightarrow \Sigma'$ such that $\sigma(e) \in \mathcal{E}'$ for every $e \in \mathcal{E}$; if, moreover, σ is a signature inclusion $\sigma: \Sigma \hookrightarrow \Sigma'$ then $\sigma: \langle \Sigma, \mathcal{E} \rangle \hookrightarrow \langle \Sigma', \mathcal{E}' \rangle$ is a *theory inclusion*. \square

Exercise 2.3.12. Let $\sigma: \langle \Sigma, \mathcal{E} \rangle \rightarrow \langle \Sigma', \mathcal{E}' \rangle$ and $\sigma': \langle \Sigma', \mathcal{E}' \rangle \rightarrow \langle \Sigma'', \mathcal{E}'' \rangle$ be theory morphisms. Show that $\sigma; \sigma': \Sigma \rightarrow \Sigma''$ is a theory morphism $\sigma; \sigma': \langle \Sigma, \mathcal{E} \rangle \rightarrow \langle \Sigma'', \mathcal{E}'' \rangle$. \square

Proposition 2.3.13. Let $\sigma: \Sigma \rightarrow \Sigma'$ be a signature morphism, \mathcal{E} be a set of Σ -equations and \mathcal{E}' be a set of Σ' -equations. Then the following conditions are equivalent:

1. σ is a theory morphism $\sigma: \langle \Sigma, Cl_{\Sigma}(\mathcal{E}) \rangle \rightarrow \langle \Sigma', Cl_{\Sigma'}(\mathcal{E}') \rangle$.

2. $\sigma(\mathcal{E}) \subseteq Cl_{\Sigma'}(\mathcal{E}')$.
3. For every $A' \in Mod_{\Sigma'}(\mathcal{E}')$, $A'|_{\sigma} \in Mod_{\Sigma}(\mathcal{E})$.

Proof. Exercise. HINT: Use the Satisfaction Lemma, Lemma 2.1.8. \square

The fact that 2.3.13(2) implies 2.3.13(1) gives a shortcut for checking if a signature morphism is a theory morphism: one need only check, for each axiom in some *presentation* of the source theory, that the translation of that axiom is in the target theory. The equivalence between 2.3.13(1) and 2.3.13(3) is similar in spirit to the Satisfaction Lemma, demonstrating a perfect correspondence between translation of syntax (axioms) along a signature morphism and translation of semantics (models) in the opposite direction. This equivalence shows that there is a model-level alternative to the axiom-level phrasing of Definition 2.3.11; in fact, we will take this alternative in the case of structured specifications (Chapter 5), where there is no equivalent axiom-level characterisation (Exercise 5.5.4).

Example 2.3.14. Let Σ be the signature

$$\begin{aligned} \Sigma = & \text{ sorts } s, BBool \\ & \text{ ops } \text{ tttt}: BBool \\ & \quad \text{ ffff}: BBool \\ & \quad \text{ not}: BBool \rightarrow BBool \\ & \quad \text{ and}: BBool \times BBool \rightarrow BBool \\ & \quad \text{ --}\leq\text{--}: s \times s \rightarrow BBool \end{aligned}$$

and recall the presentation $BOOL = \langle \Sigma_{BOOL}, \mathcal{E}_{BOOL} \rangle$ from Example 2.2.3. Define a signature morphism $\sigma: \Sigma \rightarrow \Sigma_{BOOL}$ by

$$\begin{aligned} \sigma_{\text{sorts}} &= \{s \mapsto Bool, BBool \mapsto Bool\}, \\ \sigma_{\varepsilon, BBool} &= \{\text{ tttt} \mapsto true, \text{ ffff} \mapsto false\}, \\ \sigma_{BBool, BBool} &= \{\text{ not} \mapsto \neg\}, \\ \sigma_{BBool BBool, BBool} &= \{\text{ and} \mapsto \wedge\}, \\ \sigma_{s, BBool} &= \{\leq \mapsto \Rightarrow\}. \end{aligned}$$

Let \mathcal{E} be the set of Σ -equations

$$\mathcal{E} = \{\forall x: s \bullet x \leq x = \text{ tttt}, \forall p: BBool \bullet \text{ and}(p, \text{ tttt}) = p\}.$$

Then $Cl_{\Sigma}(\mathcal{E})$ includes Σ -equations, such as $\forall p: BBool, x: s \bullet \text{ and}(p, x \leq x) = p$, that were not in \mathcal{E} . Similarly, by Example 2.3.7, $Cl_{\Sigma_{BOOL}}(\mathcal{E}_{BOOL})$ includes the Σ_{BOOL} -equation $\forall p: Bool \bullet \neg(p \wedge \neg false) = \neg p$, but it does *not* include $\forall p: Bool \bullet \neg \neg p = p$. The presentations $\langle \Sigma, Cl_{\Sigma}(\mathcal{E}) \rangle$ and $\langle \Sigma_{BOOL}, Cl_{\Sigma_{BOOL}}(\mathcal{E}_{BOOL}) \rangle$ are theories — the latter is the theory presented by $BOOL$. The signature morphism $\sigma: \Sigma \rightarrow \Sigma_{BOOL}$ is a theory morphism $\sigma: \langle \Sigma, Cl_{\Sigma}(\mathcal{E}) \rangle \rightarrow \langle \Sigma_{BOOL}, Cl_{\Sigma_{BOOL}}(\mathcal{E}_{BOOL}) \rangle$.

Recalling Example 2.2.4, the theory presented by BA is $\langle \Sigma_{BA}, Cl_{\Sigma_{BA}}(\mathcal{E}_{BA}) \rangle$, the theory of Boolean algebras, with $Cl_{\Sigma_{BA}}(\mathcal{E}_{BA})$ including, for instance, the de Morgan laws (Exercise 2.2.5). The obvious signature morphism $\iota: \Sigma_{BOOL} \rightarrow \Sigma_{BA}$ is a theory morphism $\iota: \langle \Sigma_{BOOL}, Cl_{\Sigma_{BOOL}}(\mathcal{E}_{BOOL}) \rangle \rightarrow \langle \Sigma_{BA}, Cl_{\Sigma_{BA}}(\mathcal{E}_{BA}) \rangle$.

These two theory morphisms can be composed, yielding the theory morphism $\sigma; \iota: \langle \Sigma, Cl_{\Sigma}(\mathcal{E}) \rangle \rightarrow \langle \Sigma_{BA}, Cl_{\Sigma_{BA}}(\mathcal{E}_{BA}) \rangle$. \square

Exercise 2.3.15. Give presentations $\langle \Sigma, \mathcal{E} \rangle$ and $\langle \Sigma', \mathcal{E}' \rangle$ and a theory morphism $\sigma: \langle \Sigma, Cl_\Sigma(\mathcal{E}) \rangle \rightarrow \langle \Sigma', Cl_{\Sigma'}(\mathcal{E}') \rangle$ such that $\sigma(\mathcal{E}) \not\subseteq \mathcal{E}'$. Note that this does *not* contradict the equivalence between 2.3.13(1) and 2.3.13(2). \square

2.4 Equational calculus

As we have seen, each presentation $\langle \Sigma, \mathcal{E} \rangle$ determines a theory $\langle \Sigma, Cl_\Sigma(\mathcal{E}) \rangle$, where $Cl_\Sigma(\mathcal{E})$ contains \mathcal{E} together with all of its semantic consequences. An obvious question at this point is how to determine whether or not a given Σ -equation $\forall X \bullet t = t'$ belongs to the set $Cl_\Sigma(\mathcal{E})$, i.e. how to decide if $\mathcal{E} \models_\Sigma \forall X \bullet t = t'$. The definition of $Cl_\Sigma(\mathcal{E})$ does not provide an effective method: according to this, testing $\mathcal{E} \models_\Sigma \forall X \bullet t = t'$ involves constructing the (infinite!) class $Mod_\Sigma(\mathcal{E})$ and checking whether or not $\forall X \bullet t = t'$ is satisfied by each of the algebras in this class, that is, checking for each algebra $A \in Mod_\Sigma(\mathcal{E})$ and function $v: X \rightarrow |A|$ (there may be infinitely many such functions for a given A) that $t_A(v) = t'_A(v)$. An alternative is to proceed “syntactically” by means of *inference rules* which allow the elements of $Cl_\Sigma(\mathcal{E})$ to be *derived* from the axioms in \mathcal{E} via a sequence of formal proof steps.

Throughout this section, let Σ be a signature.

Definition 2.4.1 (Equational calculus). A Σ -equation e is a *proof-theoretic consequence* of (or is *provable* from) a set \mathcal{E} of Σ -equations, written $\mathcal{E} \vdash_\Sigma e$, if this can be derived by application of the following inference rules:

$$\begin{array}{l}
 \text{(axiom)} \quad \frac{}{\mathcal{E} \vdash_\Sigma \forall X \bullet t = t'} \quad \forall X \bullet t = t' \in \mathcal{E} \\
 \\
 \text{(reflexivity)} \quad \frac{}{\mathcal{E} \vdash_\Sigma \forall X \bullet t = t} \quad X_s \subseteq \mathcal{X} \text{ for all } s \in S, \text{ and } t \in |T_\Sigma(X)| \\
 \\
 \text{(symmetry)} \quad \frac{\mathcal{E} \vdash_\Sigma \forall X \bullet t = t'}{\mathcal{E} \vdash_\Sigma \forall X \bullet t' = t} \\
 \\
 \text{(transitivity)} \quad \frac{\mathcal{E} \vdash_\Sigma \forall X \bullet t = t' \quad \mathcal{E} \vdash_\Sigma \forall X \bullet t' = t''}{\mathcal{E} \vdash_\Sigma \forall X \bullet t = t''} \\
 \\
 \text{(congruence)} \quad \frac{\mathcal{E} \vdash_\Sigma \forall X \bullet t_1 = t'_1 \quad \dots \quad \mathcal{E} \vdash_\Sigma \forall X \bullet t_n = t'_n \quad f: s_1 \times \dots \times s_n \rightarrow s \text{ in } \Sigma, \quad t_i, t'_i \in |T_\Sigma(X)|_{s_i} \text{ for } i \leq n}{\mathcal{E} \vdash_\Sigma \forall X \bullet f(t_1, \dots, t_n) = f(t'_1, \dots, t'_n)} \\
 \\
 \text{(instantiation)} \quad \frac{\mathcal{E} \vdash_\Sigma \forall X \bullet t = t'}{\mathcal{E} \vdash_\Sigma \forall Y \bullet t[\theta] = t'[\theta]} \quad \theta: X \rightarrow |T_\Sigma(Y)| \quad \square
 \end{array}$$

Exercise 2.4.2 (Admissibility of weakening and cut). Prove that if $\mathcal{E} \vdash_\Sigma \forall X \bullet t = t'$ and $\mathcal{E} \subseteq \mathcal{E}'$ then $\mathcal{E}' \vdash_\Sigma \forall X \bullet t = t'$. (HINT: Simple induction on the structure of the

derivation of $\mathcal{E} \vdash_{\Sigma} \forall X \bullet t = t'$.) This shows that the following rule is admissible²:

$$\text{(weakening)} \quad \frac{\mathcal{E} \vdash_{\Sigma} \forall X \bullet t = t'}{\mathcal{E} \cup \mathcal{E}' \vdash_{\Sigma} \forall X \bullet t = t'}$$

Prove that if $\mathcal{E} \vdash_{\Sigma} e$ and $\{e\} \cup \mathcal{E}' \vdash_{\Sigma} e'$ then $\mathcal{E} \cup \mathcal{E}' \vdash_{\Sigma} e'$. (HINT: Use induction on the structure of the derivation of $\{e\} \cup \mathcal{E}' \vdash_{\Sigma} e'$; for the case of the axiom rule, use the fact that weakening is admissible.) This shows that the following rule is admissible:

$$\text{(cut)} \quad \frac{\mathcal{E} \vdash_{\Sigma} e \quad \{e\} \cup \mathcal{E}' \vdash_{\Sigma} e'}{\mathcal{E} \cup \mathcal{E}' \vdash_{\Sigma} e'}$$

Check that your proof can be generalised to show that if $\mathcal{E} \vdash e'$ and $\mathcal{E}'_e \vdash e$ for each $e \in \mathcal{E}$ then $\bigcup_{e \in \mathcal{E}} \mathcal{E}'_e \vdash e'$. \square

Exercise 2.4.3 (Consequence is preserved by translation). Show that for any signature morphism $\sigma: \Sigma \rightarrow \Sigma'$, set \mathcal{E} of Σ -equations, and Σ -equation e , if $\mathcal{E} \vdash_{\Sigma} e$ then $\sigma(\mathcal{E}) \vdash_{\Sigma'} \sigma(e)$. \square

Example 2.4.4. Recall the presentation $\text{BOOL} = \langle \Sigma_{\text{BOOL}}, \mathcal{E}_{\text{BOOL}} \rangle$ given in Example 2.2.3. The following is a derivation of $\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p: \text{Bool} \bullet \neg(p \wedge \neg \text{false}) = \neg p$:

$$\begin{array}{c} \triangle \\ \text{P} \\ \hline \frac{\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p: \text{Bool} \bullet \neg(p \wedge \neg \text{false}) = \neg(p \wedge \text{true}) \quad \frac{\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p: \text{Bool} \bullet p \wedge \text{true} = p}{\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p: \text{Bool} \bullet \neg(p \wedge \text{true}) = \neg p}}{\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p: \text{Bool} \bullet \neg(p \wedge \neg \text{false}) = \neg p} \end{array}$$

where P is the derivation

$$\frac{\frac{\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p: \text{Bool} \bullet p = p \quad \frac{\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \neg \text{false} = \text{true}}{\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p: \text{Bool} \bullet \neg \text{false} = \text{true}}}{\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p: \text{Bool} \bullet p \wedge \neg \text{false} = p \wedge \text{true}}}{\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p: \text{Bool} \bullet \neg(p \wedge \neg \text{false}) = \neg(p \wedge \text{true})}$$

Exercise. Tag each step above with the inference rule being applied. \square

Exercise 2.4.5. Give a derivation of $\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p: \text{Bool} \bullet p \Rightarrow p = \text{true}$.

² A rule is *admissible* in a formal system of rules if its conclusion is derivable in the system provided that all its premises are derivable. This holds in particular if the rule is *derivable* in the system, that is, if it can be obtained by composition of the rules in the system.

A considerably more serious challenge is to give derivations for the de Morgan laws from the axioms of Boolean algebra (see Example 2.2.4 and Exercise 2.2.5). \square

On its own, the equational calculus is nothing more than a game with symbols; its importance lies in the correspondence between the two relations \models_{Σ} and \vdash_{Σ} . As we shall see, this correspondence is exact: \vdash_{Σ} is both *sound* and *complete* for \models_{Σ} . Soundness ($\mathcal{E} \vdash_{\Sigma} e$ implies $\mathcal{E} \models_{\Sigma} e$) is a vital property for any formal system: it ensures that the inference rules cannot be used to derive an incorrect result.

Theorem 2.4.6 (Soundness of equational calculus). *Let \mathcal{E} be a set of Σ -equations and let e be a Σ -equation. If $\mathcal{E} \vdash_{\Sigma} e$ then $\mathcal{E} \models_{\Sigma} e$.* \square

Exercise 2.4.7. Prove Theorem 2.4.6. Use induction on the depth of the derivation of $\mathcal{E} \vdash_{\Sigma} e$, showing that each rule in the system preserves the indicated property. \square

Example 2.4.8. By Theorem 2.4.6, the formal derivation in Example 2.4.4 justifies the claim in Example 2.3.7 that $\mathcal{E}_{\text{BOOL}} \models_{\Sigma_{\text{BOOL}}} \forall p:\text{Bool} \bullet \neg(p \wedge \neg\text{false}) = \neg p$. On the other hand, since $\mathcal{E}_{\text{BOOL}} \not\models_{\Sigma_{\text{BOOL}}} \forall p:\text{Bool} \bullet \neg\neg p = p$, there can be no proof in the equational calculus for $\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall p:\text{Bool} \bullet \neg\neg p = p$. \square

It is a somewhat counterintuitive fact (see [GM85]) that simplifying the calculus by omitting explicit quantifiers in equations yields an unsound system. This is due to the fact that algebras may have empty carrier sets. Any equation that includes a quantified variable $x:s$ will be satisfied by any algebra having an empty carrier for s , even if x appears on neither side of the equation. The instantiation rule is the only one that can be used to change the set of quantified variables; it is designed to ensure that quantified variables are eliminated only when it is sound to do so.

Exercise 2.4.9. Formulate a version of the equational calculus without explicit quantifiers on equations and show that it is unsound. (HINT: Consider the signature Σ with sorts s, s' and operations $f:s \rightarrow s'$, $a:s'$, $b:s'$, and set $\mathcal{E} = \{f(x) = a, f(x) = b\}$ of Σ -equations. Show that $\mathcal{E} \vdash_{\Sigma} a = b$ in your version of the calculus. Then give a Σ -algebra $A \in \text{Mod}_{\Sigma}(\mathcal{E})$ such that $A \not\models_{\Sigma} a = b$.) Pinpoint where this proof of unsoundness breaks down for the version of the equational calculus given in Definition 2.4.1. \square

Exercise 2.4.10. Show that the equational calculus without explicit quantifiers is sound when the definition of Σ -algebra is changed to require all carrier sets to be non-empty, or when either of the following constraints on Σ is imposed:

1. Σ has only one sort.
2. All sorts in Σ are *non-void*: for each sort name s in Σ , $|T_{\Sigma}|_s \neq \emptyset$. \square

Exercise 2.4.11. Give an example of a signature Σ which satisfies neither 2.4.10(1) nor 2.4.10(2) for which the equational calculus without explicit quantifiers is sound. \square

Completeness ($\mathcal{E} \models_{\Sigma} e$ implies $\mathcal{E} \vdash_{\Sigma} e$) is typically more difficult to achieve than soundness: it means that the rules in the system are powerful enough to derive all correct results. It is not as important as soundness, in the sense that a complete but unsound system is useless while (as we shall see in the sequel) a sound but incomplete system is often the best that can be obtained. The equational calculus happens to be complete for \models_{Σ} :

Theorem 2.4.12 (Completeness of equational calculus). *Let \mathcal{E} be a set of Σ -equations and let e be a Σ -equation. If $\mathcal{E} \models_{\Sigma} e$ then $\mathcal{E} \vdash_{\Sigma} e$.*

Proof sketch. Suppose $\mathcal{E} \models_{\Sigma} \forall X \bullet t = t'$. Define $\equiv \subseteq |T_{\Sigma}(X)| \times |T_{\Sigma}(X)|$ by $u \equiv u' \iff \mathcal{E} \vdash_{\Sigma} \forall X \bullet u = u'$; then \equiv is a Σ -congruence on $T_{\Sigma}(X)$. $T_{\Sigma}(X)/\equiv \models_{\Sigma} \mathcal{E}$, so $T_{\Sigma}(X)/\equiv \models_{\Sigma} \forall X \bullet t = t'$, and thus $t \equiv t'$, i.e. $\mathcal{E} \vdash_{\Sigma} \forall X \bullet t = t'$. \square

Exercise 2.4.13. Fill in the gaps in the proof of Theorem 2.4.12. \square

There are several different but equivalent versions of the equational calculus. The following exercise considers various alternatives to the congruence and instantiation rules.

Exercise 2.4.14. Show that the version of the equational calculus in Definition 2.4.1 is equivalent to the system obtained when the congruence and instantiation rules are replaced by the following single rule:

$$\text{(substitutivity)} \quad \frac{\mathcal{E} \vdash_{\Sigma} \forall X \bullet t = t' \quad \text{for each } x \in X, \mathcal{E} \vdash_{\Sigma} \forall Y \bullet \theta(x) = \theta'(x)}{\mathcal{E} \vdash_{\Sigma} \forall Y \bullet t[\theta] = t'[\theta']} \quad \theta, \theta': X \rightarrow |T_{\Sigma}(Y)|$$

Show that this is equivalent to the system having the following more restricted version of the substitutivity rule:

$$\text{(substitutivity')} \quad \frac{\mathcal{E} \vdash_{\Sigma} \forall X \cup \{x:s\} \bullet t = t' \quad \mathcal{E} \vdash_{\Sigma} \forall Y \bullet u = u'}{\mathcal{E} \vdash_{\Sigma} \forall X \cup Y \bullet t[x \mapsto u] = t'[x \mapsto u']} \quad u, u' \in |T_{\Sigma}(Y)|_s$$

(HINT: The equivalence relies on the fact that the set of quantified variables in an equation is finite.) Finally, show that both of the following rules may be derived in any of these systems:

$$\text{(abstraction)} \quad \frac{\mathcal{E} \vdash_{\Sigma} \forall X \bullet t = t'}{\mathcal{E} \vdash_{\Sigma} \forall X \cup Y \bullet t = t'} \quad Y_s \subseteq \mathcal{X} \text{ for all } s \in S$$

$$\text{(concretion)} \quad \frac{\mathcal{E} \vdash_{\Sigma} \forall X \cup \{x:s\} \bullet t = t'}{\mathcal{E} \vdash_{\Sigma} \forall X \bullet t = t'} \quad t, t' \in |T_{\Sigma}(X)| \text{ and } |T_{\Sigma}(X)|_s \neq \emptyset \quad \square$$

A consequence of the soundness and completeness theorems is that the equational calculus constitutes a *semi-decision procedure* for \models_{Σ} : enumerating all derivations will eventually produce a derivation for $\mathcal{E} \vdash_{\Sigma} e$ if $\mathcal{E} \models_{\Sigma} e$ holds, but if $\mathcal{E} \not\models_{\Sigma} e$ then this procedure will never terminate. This turns out to be the best we can achieve:

Theorem 2.4.15. *There is no decision procedure for \models_{Σ} .*

Proof. Follows immediately from the undecidability of the word problem for semi-groups [Pos47]. \square

Mechanised proof search techniques can be applied with considerable success to the discovery of derivations (and under certain conditions, discussed in Section 2.6, a decision procedure *is* possible), but Theorem 2.4.15 shows that such techniques can provide no more than a partial solution.

2.5 Initial models

The class of algebras given by the loose semantics of a Σ -presentation contains too many algebras to be very useful in practice. In particular, Birkhoff’s Variety Theorem guarantees that this class will always include degenerate Σ -algebras having a single value of each sort in Σ , as well as (nearly always) Σ -algebras that are not reachable. This unsatisfactory state of affairs is a consequence of the limited power of equational axioms. A standard way out is to take the so-called *initial semantics* of presentations, which selects a certain class of “best” models from among all those satisfying the axioms. Various alternatives to this approach will be presented in the sequel.

Throughout this section, let $\langle \Sigma, \mathcal{E} \rangle$ be a presentation.

Exercise 2.5.1. Verify the above claim concerning Birkhoff’s Variety Theorem, being specific about the meaning of “nearly always”. \square

There are two features that render certain models of presentations unfit for use in practice. The mnemonic terms “junk” and “confusion” were coined in [BG81] to characterise these:

Definition 2.5.2 (Junk and confusion). Let A be a model of $\langle \Sigma, \mathcal{E} \rangle$. We say that A *contains junk* if it is not reachable, and that A *contains confusion* if it satisfies a ground Σ -equation that is not in $Cl_{\Sigma}(\mathcal{E})$. \square

The intuition behind these terms should be readily apparent: “junk” refers to useless values which could be discarded without being missed, and “confusion” refers to the values of two ground terms being unnecessarily identified (confused).

Example 2.5.3. Recall the presentation $\text{BOOL} = \langle \Sigma_{\text{BOOL}}, \mathcal{E}_{\text{BOOL}} \rangle$ and its models $A1$, $A2$ and $A3$ from Example 2.2.3. $A1$ contains confusion ($A1 \models_{\Sigma_{\text{BOOL}}} \text{true} = \text{false} \notin Cl_{\Sigma_{\text{BOOL}}}(\mathcal{E}_{\text{BOOL}})$) but not junk; $A2$ contains junk (there is no ground Σ_{BOOL} -term t such that $t_{A2} = \spadesuit \in |A2|_{\text{Bool}}$) but not confusion; $A3$ contains neither junk nor confusion. There are models of BOOL containing both junk and confusion. (**Exercise:** Find one.) \square

Exercise 2.5.4. Consider the following specification of the natural numbers with addition:

spec NAT = **sorts** *Nat*
ops $0: \textit{Nat}$
 $\textit{succ}: \textit{Nat} \rightarrow \textit{Nat}$
 $-- + --: \textit{Nat} \times \textit{Nat} \rightarrow \textit{Nat}$
 $\forall m, n: \textit{Nat}$

- $0 + n = n$
- $\textit{succ}(m) + n = \textit{succ}(m + n)$

List some of the models of NAT. Which of these contain junk and/or confusion? (NOTE: For reference later in this section, Σ_{NAT} refers to the signature of NAT and \mathcal{E}_{NAT} refers to its axioms.) \square

Exercise 2.5.5. According to Exercise 1.3.5, surjective homomorphisms reflect junk. Show that injective homomorphisms preserve junk and reflect confusion, and that all homomorphisms preserve confusion. It follows that isomorphisms preserve and reflect junk and confusion. \square

Examples like the ones above suggest that often the algebras of interest are those which contain neither junk nor confusion. Recall Exercise 1.4.14, which characterised reachable Σ -algebras as those which are isomorphic to a quotient of T_Σ . Accordingly, the algebras we want are all isomorphic to quotients of T_Σ ; by Exercise 2.5.5 it is enough to consider just these quotient algebras themselves. Of course, not all quotients T_Σ/\equiv will be models of $\langle \Sigma, \mathcal{E} \rangle$: this will only be the case when \equiv identifies enough terms that the equations in \mathcal{E} are satisfied. But if \equiv identifies “too many” terms, T_Σ/\equiv will contain confusion. There is exactly one Σ -congruence that yields a model of $\langle \Sigma, \mathcal{E} \rangle$ containing no confusion:

Definition 2.5.6 (Congruence generated by a set of equations). The relation $\equiv_{\mathcal{E}} \subseteq |T_\Sigma| \times |T_\Sigma|$ is defined by $t \equiv_{\mathcal{E}} t' \iff \mathcal{E} \models_{\Sigma} t = t'$, for all $t, t' \in |T_\Sigma|$. $\equiv_{\mathcal{E}}$ is called the Σ -congruence generated by \mathcal{E} . \square

Exercise 2.5.7. Prove that $\equiv_{\mathcal{E}}$ is a Σ -congruence on T_Σ . \square

Theorem 2.5.8 (Quotient construction). $T_\Sigma/\equiv_{\mathcal{E}}$ is a model of $\langle \Sigma, \mathcal{E} \rangle$ containing no junk and no confusion. \square

Exercise 2.5.9. Prove Theorem 2.5.8. HINT: Note that $T_\Sigma/\equiv_{\mathcal{E}}$ contains no junk by Exercise 1.4.14. Then show that for any term $t \in T_\Sigma(X)$ and substitution $\theta: X \rightarrow T_\Sigma$, $t_{T_\Sigma/\equiv_{\mathcal{E}}}(\theta') = [t[\theta]]_{\equiv_{\mathcal{E}}}$, where $\theta'(x) = [\theta(x)]_{\equiv_{\mathcal{E}}}$ for $x \in X$. Use this to show that $T_\Sigma/\equiv_{\mathcal{E}}$ satisfies all the equations in \mathcal{E} and contains no confusion. \square

Example 2.5.10. Recall again the presentation $\text{BOOL} = \langle \Sigma_{\text{BOOL}}, \mathcal{E}_{\text{BOOL}} \rangle$ from Example 2.2.3. The model $T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}$ of BOOL is defined as follows:

$$\begin{aligned}
|T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}|_{\text{Bool}} &= \{[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}, [false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}\}, \\
true_{T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}} &= [true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}, \\
false_{T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}} &= [false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}, \\
\neg_{T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}} &= \{[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} \mapsto [false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}, [false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} \mapsto [true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}\},
\end{aligned}$$

$\wedge_{T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$
$[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$
$[false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$

$\Rightarrow_{T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$
$[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$
$[false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$	$[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$

where

$$\begin{aligned}
[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} &= \{true, \neg false, true \wedge true, \neg(false \wedge true), \\
&\quad \neg(false \wedge \neg false), false \Rightarrow false, \dots\}, \\
[false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} &= \{false, \neg true, true \wedge false, \neg(true \wedge true), \\
&\quad \neg(true \wedge \neg false), true \Rightarrow false, \dots\}.
\end{aligned}$$

The carrier set $|T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}|_{\text{Bool}}$ has just two elements since the axioms in $\mathcal{E}_{\text{BOOL}}$ can be used to reduce each ground Σ_{BOOL} -term to *true* or *false*, and $true \neq_{\mathcal{E}_{\text{BOOL}}} false$. Note that the “syntactic” nature of $T_{\Sigma_{\text{BOOL}}}$ is preserved in $T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}$, e.g. for each $x \in [true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$, “ $\neg x$ ” $\in [false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} = \neg_{T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}}([true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}})$. \square

Exercise 2.5.11. Recall the presentation $\text{NAT} = \langle \Sigma_{\text{NAT}}, \mathcal{E}_{\text{NAT}} \rangle$ from Exercise 2.5.4. Construct the model $T_{\Sigma_{\text{NAT}}}/\equiv_{\mathcal{E}_{\text{NAT}}}$ of NAT. \square

Exercise 2.5.12. Show that $\equiv_{\mathcal{E}}$ is the only Σ -congruence making Theorem 2.5.8 hold. \square

The special properties of $T_{\Sigma}/\equiv_{\mathcal{E}}$ described by Theorem 2.5.8 can be captured very succinctly by saying that $T_{\Sigma}/\equiv_{\mathcal{E}}$ is a so-called *initial model* of $\langle \Sigma, \mathcal{E} \rangle$.

Definition 2.5.13 (Initial model of a presentation). A Σ -algebra A is *initial* in a class \mathcal{A} of Σ -algebras if $A \in \mathcal{A}$ and for every $B \in \mathcal{A}$ there is a unique Σ -homomorphism $h: A \rightarrow B$. An *initial model* of $\langle \Sigma, \mathcal{E} \rangle$ is a Σ -algebra that is initial in $\text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$. $\text{IMod}[\langle \Sigma, \mathcal{E} \rangle]$ is the class of all initial models of $\langle \Sigma, \mathcal{E} \rangle$. \square

In the next chapter we will see that this definition can be generalised to a much wider context than that of algebras and homomorphisms.

Theorem 2.5.14 (Initial model theorem). $T_{\Sigma}/\equiv_{\mathcal{E}}$ is an initial model of $\langle \Sigma, \mathcal{E} \rangle$.

Proof sketch. $T_\Sigma/\equiv_{\mathcal{E}}$ is a model of $\langle \Sigma, \mathcal{E} \rangle$ by Theorem 2.5.8. For $B \in \text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$, let $\varnothing^\sharp: T_\Sigma \rightarrow B$ be the unique homomorphism from the algebra of ground Σ -terms to B . Since $B \models_\Sigma \mathcal{E}$, we have $\equiv_{\mathcal{E}} \subseteq \ker(\varnothing^\sharp)$, and by Exercise 1.3.20 there is a homomorphism $h: T_\Sigma/\equiv_{\mathcal{E}} \rightarrow B$ which is unique by Exercise 1.3.6. (**Exercise:** Fill in the gaps in this proof.) \square

Example 2.5.15. Recall the presentation $\text{BOOL} = \langle \Sigma_{\text{BOOL}}, \mathcal{E}_{\text{BOOL}} \rangle$ and its models $A1$, $A2$ and $A3$ from Example 2.2.3, and its model $T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}$ from Example 2.5.10, which is an initial model by Theorem 2.5.14. Σ_{BOOL} -homomorphisms from $T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}$ to $A1$, $A2$ and $A3$ are as follows:

$$\begin{aligned} h1: T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}} &\rightarrow A1 & h1_{\text{Bool}} &= \{[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} \mapsto \star, [false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} \mapsto \star\}, \\ h2: T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}} &\rightarrow A2 & h2_{\text{Bool}} &= \{[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} \mapsto \clubsuit, [false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} \mapsto \heartsuit\}, \\ h3: T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}} &\rightarrow A3 & h3_{\text{Bool}} &= \{[true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} \mapsto 1, [false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}} \mapsto 0\}. \end{aligned}$$

(**Exercise:** Check uniqueness.)

$A1$ is not an initial model: for example, there is no homomorphism from $A1$ to $A2$, nor from $A1$ to $A3$. In general, models containing confusion cannot be initial since homomorphisms preserve confusion (Exercise 2.5.5). Similarly, $A2$ is not an initial model: for example, there is no homomorphism from $A2$ to $A3$, since there is no value in $|A3|_{\text{Bool}}$ to which such a homomorphism could map the “extra” value $\spadesuit \in |A2|_{\text{Bool}}$. On the other hand, $A3$ is initial: for example, there is a unique homomorphism $g1: A3 \rightarrow A1$ (where $g1_{\text{Bool}}(1) = g1_{\text{Bool}}(0) = \star$), there is a unique homomorphism $g2: A3 \rightarrow A2$ (where $g2_{\text{Bool}}(1) = \clubsuit$ and $g2_{\text{Bool}}(0) = \heartsuit$), and there is a unique homomorphism $g: A3 \rightarrow T_{\Sigma_{\text{BOOL}}}/\equiv_{\mathcal{E}_{\text{BOOL}}}$ (where $g_{\text{Bool}}(1) = [true]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$ and $g_{\text{Bool}}(0) = [false]_{\equiv_{\mathcal{E}_{\text{BOOL}}}}$). \square

Exercise 2.5.16. Recall the model you constructed in Exercise 2.5.11 of the specification NAT of natural numbers with addition. Show that there is a unique homomorphism from this model to each of the models you considered in Exercise 2.5.4. \square

Exercise 2.5.17. Using Theorem 2.5.14, show that T_Σ is an initial model of $\langle \Sigma, \emptyset \rangle$. Contemplate how this relates to Fact 1.4.4 and Definition 1.4.5. \square

Exercise 2.5.18. Note that initial models of $\langle \Sigma, \mathcal{E} \rangle$ may have empty carriers for some sorts. Show that this is necessary: give an example of a presentation $\langle \Sigma, \mathcal{E} \rangle$ such that no algebra is initial in the class of its models that have non-empty carriers of all sorts. Link this with Exercise 1.2.3. \square

Taking a presentation $\langle \Sigma, \mathcal{E} \rangle$ to denote the class $\text{IMod}[\langle \Sigma, \mathcal{E} \rangle]$ of its initial models is called taking its *initial semantics*. We know from Theorem 2.5.14 that $\text{IMod}[\langle \Sigma, \mathcal{E} \rangle]$ is never empty. Although the motivation for wishing to exclude models containing junk and confusion was merely to weed out certain kinds of degenerate cases, the effect of this constraint is to restrict attention to an isomorphism class of models:

Exercise 2.5.19. Show that any two initial models of a presentation are isomorphic. Conclude that the initial models of a presentation are exactly those containing no junk and no confusion. \square

For some purposes, restricting attention to an isomorphism class of models is clearly inappropriate. The following exercise demonstrates what can go wrong.

Exercise 2.5.20. Consider the addition of a subtraction operation $-- -: Nat \times Nat \rightarrow Nat$ to the specification NAT in Exercise 2.5.4, with the axioms $\forall m: Nat \bullet m - 0 = m$ and $\forall m, n: Nat \bullet succ(m) - succ(n) = m - n$. These axioms do not fix the value of $m - n$ when $n > m$; assume that we are willing to accept any value in this case, perhaps because we are certain for some reason that it will never arise. Construct an initial model of this specification. Why is this model unsatisfactory? Can you think of a better model? What is the problem with restricting to an isomorphism class of models of this specification? \square

The phenomenon illustrated here arises in cases where operations are not defined in a *sufficiently complete* way. Roughly speaking, a definition of an operation is sufficiently complete when the value produced by the operation is defined for all of the possible values of its arguments. See Definition 6.1.22 below for a proper formulation of this property in a more general context.

One may argue that Exercise 2.5.20 is unconvincing, since the lack of sufficient completeness arises there because we do not really need $m - n$ to be defined as a natural number when $n > m$, and that this can be dealt with using one of the approaches to partial functions below (Sections 2.7.3, 2.7.4, or 2.7.5). However, the same phenomenon arises in other cases as well:

Exercise 2.5.21. Give a specification of natural numbers with a function that for each natural number n chooses an arbitrary number that is greater than n . HINT: You may first extend the specification NAT of Exercise 2.5.4 with a sort *Bool* with operations and axioms as in BOOL in Example 2.2.3, and add a binary operation $-- < -: Nat \times Nat \rightarrow Bool$ with the following axioms:

$$\begin{aligned} \forall n: Nat \bullet 0 < succ(n) &= true \\ \forall m: Nat \bullet succ(m) < 0 &= false \\ \forall m, n: Nat \bullet succ(m) < succ(n) &= m < n \end{aligned}$$

The required function $ch: Nat \rightarrow Nat$ may now be constrained by the obvious axiom $\forall n: Nat \bullet n < ch(n) = true$.

Clearly, the definition of ch cannot be sufficiently complete. Construct the initial model of the resulting specification and check that it is not satisfactory. Referring to other algebraic approaches presented in Sections 2.7.3, 2.7.4, and 2.7.5 below, check that none of them offers a satisfactory solution either. \square

The above exercise indicates one of the most compelling reasons for considering alternatives to initial semantics: requiring specifications to define all operations in a sufficiently complete way is much too restrictive in many practical cases. Such a requirement is also undesirable for methodological reasons, since it forces the specifier of a problem to make decisions which are more appropriately left to the implementor.

The comments above notwithstanding, there are certain common situations in which initial semantics is appropriate and useful. In particular, the implicit “no junk”

constraint conveniently captures the “that’s all there is” condition which is needed in inductive definitions of syntax.

Example 2.5.22. Consider the following specification of syntax for simple arithmetic expressions:

spec $\text{EXPR} = \text{sorts } \text{Expr}$
ops $x: \text{Expr}$
 $y: \text{Expr}$
 $0: \text{Expr}$
 $\text{plus}: \text{Expr} \times \text{Expr} \rightarrow \text{Expr}$
 $\text{minus}: \text{Expr} \times \text{Expr} \rightarrow \text{Expr}$
 $\forall e, e': \text{Expr}$
 $\bullet \text{plus}(e, e') = \text{plus}(e', e)$

The axiom requires the *syntax* of addition to be commutative. In the initial semantics of EXPR , the “no junk” condition ensures that the only expressions (values of sort Expr) are those built from 0 , x and y using plus and minus . The “no confusion” condition ensures that no undesired identification of expressions occurs: for example, the syntax of addition is not associative and the syntax of subtraction is not commutative. \square

Exercise 2.5.23. Write a specification of (finite) sets of natural numbers. The operations should include $\emptyset: \text{NatSet}$, $\text{singleton}: \text{Nat} \rightarrow \text{NatSet}$ and $_ \cup _: \text{NatSet} \times \text{NatSet} \rightarrow \text{NatSet}$. \square

The “no junk” condition is more powerful than it might appear to be at first glance. Imposing the constraint that every value be expressible as a ground term makes it possible to use induction on the structure of terms to prove properties of all the values in an algebra. This means that for reasoning about models of specifications containing no junk, such as initial models, it is sound to add an induction rule scheme to the equational calculus presented in the previous section. Since the form of the induction rule scheme varies according to the signature of the specification at hand, this is best illustrated by means of examples.

Example 2.5.24. Recall the presentation $\text{NAT} = \langle \Sigma_{\text{NAT}}, \mathcal{E}_{\text{NAT}} \rangle$ of natural numbers with addition given in Exercise 2.5.4. To simplify notation, let x and y stand for variable names such that $x: \text{Nat}$ and $y: \text{Nat}$ are not in Σ_{NAT} and $x: \text{Nat}$ does not appear in the $\text{sorts}(\Sigma_{\text{NAT}})$ -sorted set of variables X used below. The following induction rule scheme is sound for reachable models of NAT (and for reachable models of all other Σ_{NAT} -presentations):

$$\frac{\mathcal{E} \vdash_{\Sigma_{\text{NAT}}} P(0) \quad \mathcal{E} \cup \{P(x)\} \vdash_{\Sigma_{\text{NAT}} \cup \{x: \text{Nat}\}} P(\text{succ}(x)) \quad \mathcal{E} \cup \{P(x), P(y)\} \vdash_{\Sigma_{\text{NAT}} \cup \{x, y: \text{Nat}\}} P(x + y)}{\mathcal{E} \vdash_{\Sigma_{\text{NAT}}} \forall x: \text{Nat} \bullet P(x)}$$

Here, $P(x)$ stands for a $(\Sigma_{\text{NAT}} \cup \{x: \text{Nat}\})$ -equation, $\forall X \bullet t = t'$; think of this as a Σ_{NAT} -equation with free variable $x: \text{Nat}$. Then $P(0)$ stands for the Σ_{NAT} -equation

$\forall X \bullet t[x \mapsto 0] = t'[x \mapsto 0]$, $P(\text{succ}(x))$ stands for the $(\Sigma_{\text{NAT}} \cup \{x:\text{Nat}\})$ -equation $\forall X \bullet t[x \mapsto \text{succ}(x)] = t'[x \mapsto \text{succ}(x)]$, and analogously for $P(y)$ and $P(x+y)$, and $\forall x:\text{Nat} \bullet P(x)$ stands for the Σ_{NAT} -equation $\forall X \cup \{x:\text{Nat}\} \bullet t = t'$. The following additional inference rule is needed to infer equations over $\Sigma_{\text{NAT}} \cup \{x:\text{Nat}\}$ and $\Sigma_{\text{NAT}} \cup \{x, y:\text{Nat}\}$ from Σ_{NAT} -equations:

$$\frac{\mathcal{E} \vdash_{\Sigma} \forall X \bullet t = t'}{\mathcal{E} \vdash_{\Sigma \cup \Sigma'} \forall X \bullet t = t'}$$

Exercise. Show that adding the two inference rules above to the equational calculus gives a system that is sound for reachable models of Σ_{NAT} -presentations.

The inference rule scheme above can be used for proving theorems such as associativity and commutativity of $+$. But note that the axioms for $+$ fully define it in terms of 0 and succ : it is possible to prove by induction on the structure of terms that for every ground Σ_{NAT} -term t there is a ground Σ_{NAT} -term t' such that t' does not contain the $+$ operation and $\mathcal{E}_{\text{NAT}} \vdash_{\Sigma_{\text{NAT}}} t = t'$. (**Exercise:** Prove it. Note that this is a proof at the meta-level *about* \vdash , not a derivation at the object level *using* \vdash .) This shows that the third premise of the above induction rule scheme is redundant. Eliminating it gives the following scheme, which is more obviously related to the usual form of induction for natural numbers:

$$\frac{\mathcal{E} \cup \mathcal{E}_{\text{NAT}} \vdash_{\Sigma_{\text{NAT}}} P(0) \quad \mathcal{E} \cup \mathcal{E}_{\text{NAT}} \cup \{P(x)\} \vdash_{\Sigma_{\text{NAT}} \cup \{x:\text{Nat}\}} P(\text{succ}(x))}{\mathcal{E} \cup \mathcal{E}_{\text{NAT}} \vdash_{\Sigma_{\text{NAT}}} \forall x:\text{Nat} \bullet P(x)}$$

Taking $P(x)$ to be $\forall n, p:\text{Nat} \bullet x + (n+p) = (x+n) + p$, we have the following derivation, which proves that addition is associative in initial models of NAT (**Exercise:** Supply the derivations P_1 and P_2):

$$\frac{\begin{array}{c} \text{---} \\ \diagdown \quad \diagup \\ P_1 \\ \diagup \quad \diagdown \\ \text{---} \\ \mathcal{E}_{\text{NAT}} \vdash_{\Sigma_{\text{NAT}}} \forall n, p:\text{Nat} \bullet \\ 0 + (n+p) = (0+n) + p \end{array} \quad \begin{array}{c} \text{---} \\ \diagdown \quad \diagup \\ P_2 \\ \diagup \quad \diagdown \\ \text{---} \\ \mathcal{E}_{\text{NAT}} \cup \{\forall n, p:\text{Nat} \bullet x + (n+p) = (x+n) + p\} \\ \vdash_{\Sigma_{\text{NAT}} \cup \{x:\text{Nat}\}} \forall n, p:\text{Nat} \bullet \\ \text{succ}(x) + (n+p) = (\text{succ}(x) + n) + p \end{array}}{\mathcal{E}_{\text{NAT}} \vdash_{\Sigma_{\text{NAT}}} \forall x, n, p:\text{Nat} \bullet x + (n+p) = (x+n) + p}$$

There are models of NAT containing junk which do not satisfy $\forall x, n, p:\text{Nat} \bullet x + (n+p) = (x+n) + p$. Hence, this equation is not in $Cl_{\Sigma_{\text{NAT}}}(\mathcal{E}_{\text{NAT}})$ and induction is required for its derivation. \square

Exercise 2.5.25. Recall again the presentation $\text{BOOL} = \langle \Sigma_{\text{BOOL}}, \mathcal{E}_{\text{BOOL}} \rangle$ from Example 2.2.3. Give an induction rule scheme that is sound for reachable models of Σ_{BOOL} -presentations. (HINT: There will be five premises, one for each operation in BOOL.) Show that three of the premises are redundant (HINT: eliminate one operation at a time), which gives the following rule scheme:

$$\frac{\mathcal{E} \cup \mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} P(\text{true}) \quad \mathcal{E} \cup \mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} P(\text{false})}{\mathcal{E} \cup \mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \forall x:\text{Bool} \bullet P(x)}$$

Use this to prove that $\forall p:\text{Bool} \bullet \neg\neg p = p$ holds in initial models of **BOOL**. Prove that the axiom $\forall p:\text{Bool} \bullet p \wedge \neg p = \text{false}$ is redundant for the initial semantics of **BOOL**, that is:

$$\mathcal{E}_{\text{BOOL}} \setminus \{\forall p:\text{Bool} \bullet p \wedge \neg p = \text{false}\} \vdash_{\Sigma_{\text{BOOL}}} \forall p:\text{Bool} \bullet p \wedge \neg p = \text{false}. \quad \square$$

Adding an induction rule scheme appropriate to the signature at hand to the equational calculus gives a system that is sound for reasoning about initial models of specifications, and is more powerful than the equational calculus on its own. However, the resulting system is not always complete. In fact, it turns out that completeness is unachievable in general: there is *no* sound proof system that is complete for reasoning about initial models of arbitrary specifications. In order to prove that this is the case, it is necessary to formalise what we mean by the term “proof system”. For our purposes it will suffice to assume that any proof system has a recursively enumerable set of theorems. See [Chu56] for a discussion of the philosophical considerations (e.g. finiteness of proofs, decidability of the correctness of individual proof steps) underlying this assumption.

Theorem 2.5.26 (Incompleteness for initial semantics). *There is a presentation $\langle \Sigma, \mathcal{E} \rangle$ such that there is no proof system which is sound and complete with respect to satisfaction of equations in the class of initial models of $\langle \Sigma, \mathcal{E} \rangle$.*

Proof. As a consequence of Matiyasevich’s theorem, the set of equations which hold in the standard model of the natural numbers (with 0, *succ*, +, \times and $-$, such that $m - n = 0$ when $n \geq m$) is not recursively enumerable [DMR76, Sect. 8]. Therefore, this cannot be the set of theorems produced by any proof system. It is easy to construct a (single-sorted) presentation having this as an initial model. (**Exercise:** Construct it.) Since all the initial models of a presentation are isomorphic (Exercise 2.5.19) and since isomorphisms preserve and reflect satisfaction of equations (Exercise 2.1.5), this completes the proof. \square

The fact that completeness cannot be achieved is of no real importance in practice: the equational calculus together with induction is perfectly adequate for normal use. But the failure of completeness does mean that care must be taken to distinguish between semantic consequence (\models) and provability (\vdash) in theoretical work. It is important to recognize that semantic consequence is the relation of primary importance, since it is based directly on satisfaction, which embodies *truth*. Provability is merely an *approximation* to truth, albeit one that is of great importance for practical use since it is based on mechanical syntactic manipulation. The failure of completeness means that the approximation cannot be exact, but by being sound it errs on the side of safety.

Exercise 2.5.27. Show that the equational calculus (without added induction rule schemes) is complete with respect to satisfaction of *ground* equations in initial models of specifications. \square

The additional specification techniques introduced in Chapter 5 will lead to a widening of the gap between satisfaction and provability. In particular, even completeness with respect to satisfaction of ground equations will be impossible to retain.

A generalisation of the concept of initial model is needed to give a fully satisfactory specification of classes of models that are naturally parametric with respect to some basic data. An example is the definition of terms in Section 1.4, which is parametric in an S -sorted set of variables. Another is the specification of sets (see Exercise 2.5.23): it should be possible to specify sets without building in a specification of the kind of values in the sets (in this case, natural numbers).

Exercise 2.5.28. Suppose that all information about the natural numbers is removed from the specification of sets you gave in Exercise 2.5.23, by deleting operations on natural numbers like *succ* and changing the sort name *Nat* to *Elem*. Construct an initial model of the resulting specification. Why is this model unsatisfactory? \square

The required concept is that of a *free* model extending a given algebra, which captures the idea of initiality *relative to* a fixed part of the model. See Section 3.5 for the details, Section 4.3 for the use of this concept in the context of specifications, and Chapter 6 for much more on the general topic of parameterisation.

2.6 Term rewriting

Although there is no decision procedure for \models_{Σ} (Theorem 2.4.15), there is a class of specifications for which consequence can be decided. The idea is similar to the one behind the strategy used in mathematics for proving that an equation follows from a set of equational axioms: one applies the axioms in an attempt to reduce both sides of the equation to a common result, and if this is successful then the equation follows from the axioms. An essential ingredient of this strategy is the use of equations as directed *simplification* or *rewrite rules*.

Throughout this section, let $\Sigma = \langle S, \Omega \rangle$ be a signature, and let X be an S -sorted set of variables such that $X_s \subseteq \mathcal{X}$ for all $s \in S$.

Assumption. For simplicity of presentation, we assume throughout this section that either Σ has only one sort, or all sorts in Σ are non-void (see Exercise 2.4.10). Under this assumption, the version of the equational calculus without explicit quantifiers is sound, and all references to the calculus below are to this version. See Exercises 2.6.11 and 2.6.26 for hints on how to do away with this assumption. \square

Definition 2.6.1 (Context). A Σ -context for sort $s \in S$ is a term $C \in |T_{\Sigma}(X \uplus \{\square : s\})|$ containing one occurrence of the distinguished variable \square . We write $C[\square]$ to suggest that C should be viewed as a term with a hole in it. Substitution of a term $t \in |T_{\Sigma}(X)|_s$ in $C[\square]$ gives the term $C[\square : s \mapsto t] \in |T_{\Sigma}(X)|$, written $C[t]$. \square

Definition 2.6.2 (Rewrite rule). A Σ -rewrite rule r of sort $s \in S$ consists of two Σ -terms $t, t' \in |T_{\Sigma}(X)|_s$, written $t \rightarrow t'$. A Σ -rewrite rule $r = t \rightarrow t'$ of sort s determines

a set of *reduction steps* $C[t[\theta]] \rightarrow_r C[t'[\theta]]$ for all Σ -contexts $C[\square]$ for sort s and substitutions $\theta: X \rightarrow |T_\Sigma(X)|$; this defines the relation $\rightarrow_r \subseteq |T_\Sigma(X)| \times |T_\Sigma(X)|$, the *one-step reduction relation generated by r* . The inverse of one-step reduction \rightarrow_r is *one-step expansion*, written $r \leftarrow$. \square

A reduction step $u \rightarrow_r u'$ according to a rewrite rule $r = t \rightarrow t'$ is an application of an *instance* $t[\theta] \rightarrow t'[\theta]$ of r to replace the *subterm* $t[\theta]$ of u (corresponding to the “hole” in $C[\square]$) by $t'[\theta]$. The subterm $t[\theta]$ of u is called a *redex* (short for “reducible expression”).

Definition 2.6.3 (Term rewriting system). A Σ -term rewriting system R is a set of Σ -rewrite rules. The *set of Σ -equations determined by R* is $Eq(R) = \{t = t' \mid t \rightarrow t' \in R\}$; by the assumption, we can dispense with explicit quantification of variables in equations. The *one-step reduction relation generated by R* is the relation

$$\rightarrow_R = \bigcup_{r \in R} \rightarrow_r \quad (\subseteq |T_\Sigma(X)| \times |T_\Sigma(X)|).$$

The inverse of one-step reduction \rightarrow_R is *one-step expansion*, written $R \leftarrow$. \square

Given a set \mathcal{E} of Σ -equations, a Σ -term rewriting system R will be of greatest relevance to \mathcal{E} when $Cl_\Sigma(\mathcal{E}) = Cl_\Sigma(Eq(R))$. One way to obtain such an R is to use the equations themselves as rewrite rules by selecting an *orientation* for each equation $t = t'$: either $t \rightarrow t'$ or $t' \rightarrow t$. For reasons that will become clear below, the most useful orientation is the one in which the right-hand side of the rule is “simpler” than the left-hand side. It is not always obvious how to measure simplicity of terms — in fact, this is a major issue in the theory of term rewriting — and sometimes there is no satisfactory orientation, as in the case of an equation such as $n + m = m + n$.

In the rest of this section, let R be a Σ -term rewriting system.

Definition 2.6.4 (Reduction \rightarrow_R^* and convertibility \sim_R). The *reduction relation* $\rightarrow_R^* \subseteq |T_\Sigma(X)| \times |T_\Sigma(X)|$ generated by R is the transitive reflexive closure of \rightarrow_R . In other words, $t \rightarrow_R^* t'$ if $t = t'$ or there exist terms $t_1, \dots, t_n \in |T_\Sigma(X)|$, $n \geq 0$, such that $t \rightarrow_R t_1 \rightarrow_R \dots \rightarrow_R t_n \rightarrow_R t'$; then we say that t *reduces to* t' . The inverse of reduction \rightarrow_R^* is *expansion*, written $R \leftarrow^*$. The *convertibility relation* $\sim_R \subseteq |T_\Sigma(X)| \times |T_\Sigma(X)|$ generated by R is the symmetric transitive reflexive closure of \rightarrow_R . In other words, $t \sim_R t'$ if $t = t'$ or there exist terms $t_1, \dots, t_n \in |T_\Sigma(X)|$, $n \geq 0$, such that $t \rightarrow_R t_1$ or $t_1 R \leftarrow t$, and $t_1 \rightarrow_R t_2$ or $t_1 R \leftarrow t_2$, and \dots , and $t_n \rightarrow_R t'$ or $t_n R \leftarrow t'$; then we say that t *converts to* t' . \square

Exercise 2.6.5. Check that \sim_R is a Σ -congruence on $T_\Sigma(X)$. \square

Example 2.6.6. Recall again the presentation $\text{BOOL} = \langle \Sigma_{\text{BOOL}}, \mathcal{E}_{\text{BOOL}} \rangle$ from Example 2.2.3. The following Σ_{BOOL} -term rewriting system RBOOL obviously satisfies $Cl_{\Sigma_{\text{BOOL}}}(\mathcal{E}_{\text{BOOL}}) = Cl_{\Sigma_{\text{BOOL}}}(Eq(\text{RBOOL}))$:

$$\text{RBOOL} = \{ \neg \text{true} \rightarrow \text{false}, \neg \text{false} \rightarrow \text{true}, p \wedge \text{true} \rightarrow p, p \wedge \text{false} \rightarrow \text{false}, \\ p \wedge \neg p \rightarrow \text{false}, p \Rightarrow q \rightarrow \neg(p \wedge \neg q) \}.$$

(Observe that in the rule $p \Rightarrow q \rightarrow \neg(p \wedge \neg q)$, the right-hand side is not obviously simpler than the left-hand side.) For example, we have (the redex reduced by each step is underlined>)

$$\begin{aligned}
\neg(p \wedge \underline{(q \Rightarrow \neg \text{false})}) &\rightarrow_{\text{RBooL}} \neg(p \wedge \neg(q \wedge \underline{\neg \text{false}})) \\
&\rightarrow_{\text{RBooL}} \neg(p \wedge \neg(q \wedge \underline{\neg \text{true}})) \\
&\rightarrow_{\text{RBooL}} \neg(p \wedge \neg(q \wedge \underline{\text{false}})) \\
&\rightarrow_{\text{RBooL}} \neg(p \wedge \underline{\neg \text{false}}) \\
&\rightarrow_{\text{RBooL}} \neg(\underline{p \wedge \text{true}}) \\
&\rightarrow_{\text{RBooL}} \neg p
\end{aligned}$$

so $\neg(p \wedge (q \Rightarrow \neg \text{false})) \rightarrow_{\text{RBooL}}^* \neg p$, and

$$\begin{aligned}
\neg(p \wedge (q \Rightarrow \text{false})) &\text{RBooL} \leftarrow \neg(p \wedge (q \Rightarrow \neg \text{true})) \\
&\rightarrow_{\text{RBooL}} \neg(p \wedge \neg(q \wedge \underline{\neg \text{true}})) \\
&\rightarrow_{\text{RBooL}} \neg(p \wedge \neg(q \wedge \underline{\text{false}})) \\
\text{RBooL} \leftarrow &\neg(p \wedge \neg((q \wedge \text{true}) \wedge \underline{\neg \text{false}})) \\
&\rightarrow_{\text{RBooL}} \neg(p \wedge \neg((q \wedge \text{true}) \wedge \underline{\text{true}})) \\
&\rightarrow_{\text{RBooL}} \neg(p \wedge \neg(q \wedge \underline{\text{true}}))
\end{aligned}$$

so $\neg(p \wedge (q \Rightarrow \text{false})) \sim_{\text{RBooL}} \neg(p \wedge \neg(q \wedge \text{true}))$. \square

Exercise 2.6.7. Recall the presentation $\text{NAT} = \langle \Sigma_{\text{NAT}}, \mathcal{E}_{\text{NAT}} \rangle$ from Exercise 2.5.4. Looking at the equations in \mathcal{E}_{NAT} , give a Σ_{NAT} -term rewriting system RNAT such that $Cl_{\Sigma_{\text{NAT}}}(\mathcal{E}_{\text{NAT}}) = Cl_{\Sigma_{\text{NAT}}}(Eq(\text{RNAT}))$, and practice reducing and converting some Σ_{NAT} -terms using RNAT . \square

The convertibility relation generated by R coincides with equality provable from $Eq(R)$. This fact is captured by the following two theorems.

Theorem 2.6.8 (Soundness of convertibility). *If $t \sim_R t'$ then $Eq(R) \vdash_{\Sigma} t = t'$.*

Proof sketch. Consider a reduction step $C[t[\theta]] \rightarrow_R C[t'[\theta]]$. This corresponds to a derivation involving an application of the axiom rule, to derive $Eq(R) \vdash t = t'$; an application of instantiation, to derive $Eq(R) \vdash t[\theta] = t'[\theta]$; and repeated applications of reflexivity and congruence, to derive $Eq(R) \vdash C[t[\theta]] = C[t'[\theta]]$. The definition of \sim_R as the symmetric transitive reflexive closure of \rightarrow_R corresponds directly to applications of the symmetry, transitivity and reflexivity rules. (**Exercise:** Fill in the gaps in this proof.) \square

Lemma 2.6.9. *Suppose $t, t' \in |T_{\Sigma}(X)|_s$ for $s \in S$. If $t \sim_R t'$ then*

1. $C[t] \sim_R C[t']$ for any Σ -context $C[\square]$ for sort s .
2. $t[\theta] \sim_R t'[\theta]$ for any substitution $\theta: X \rightarrow |T_{\Sigma}(X)|$.

Proof. Exercise: Do it. \square

Theorem 2.6.10 (Completeness of convertibility). *If $Eq(R) \vdash_{\Sigma} t = t'$ then $t \sim_R t'$.*

Proof sketch. By induction on the depth of the derivation of $Eq(R) \vdash_{\Sigma} t = t'$. The most interesting case is when the last step is an application of the congruence rule:

$$\frac{Eq(R) \vdash_{\Sigma} t_1 = t'_1 \quad \cdots \quad Eq(R) \vdash_{\Sigma} t_n = t'_n}{Eq(R) \vdash_{\Sigma} f(t_1, \dots, t_n) = f(t'_1, \dots, t'_n)}$$

where $f: s_1 \times \cdots \times s_n \rightarrow s$. By the inductive assumption, $t_1 \sim_R t'_1$ and \dots and $t_n \sim_R t'_n$. Then, by repeated application of Lemma 2.6.9(1), we have $f(t_1, t_2, \dots, t_n) \sim_R f(t'_1, t_2, \dots, t_n) \sim_R \cdots \sim_R f(t'_1, t'_2, \dots, t'_n)$ (using first the context $f(\square: s_1, t_2, \dots, t_n)$, then $f(t'_1, \square: s_2, \dots, t_n)$, then \dots , then $f(t'_1, t'_2, \dots, \square: s_n)$). When the last step of the derivation of $Eq(R) \vdash_{\Sigma} t = t'$ is an application of the instantiation rule, the result follows directly by Lemma 2.6.9(2). (**Exercise:** Complete the proof.) \square

Exercise 2.6.11. Try to get rid of the need for the assumption on Σ made at the beginning of this section in all the definitions and results above. This will involve rewriting terms of the form $(X)t$ using rewrite rules of the form $\forall X \bullet t \rightarrow t'$, in both cases with explicit variable declarations. \square

Given the exact correspondence between convertibility and provable equality, a decision procedure for $t \sim_R t'$ amounts to a decision procedure for $\mathcal{E} \vdash_{\Sigma} t = t'$, provided $Cl_{\Sigma}(\mathcal{E}) = Cl_{\Sigma}(Eq(R))$. The problem with testing $t \sim_R t'$ by simply applying the definition is that the “path” from t to t' may include both reduction steps and expansion steps, and may be of arbitrary length. But when R satisfies certain conditions, it is sufficient to test just a *single* path having the special form $t \rightarrow_R^* t'' \leftarrow_R^* t'$, which yields a simple and efficient decision procedure for convertibility.

Definition 2.6.12 (Normal form). A Σ -term $t \in T_{\Sigma}(X)$ is a *normal form* (for R) if there is no term t' such that $t \rightarrow_R t'$. \square

Definition 2.6.13 (Termination). A Σ -term rewriting system R is *terminating* (or *strongly normalising*) if there is no infinite reduction sequence $t_1 \rightarrow_R t_2 \rightarrow_R \cdots$; that is, whenever $t_1 \rightarrow_R t_2 \rightarrow_R \cdots$, there is some (finite) $n \geq 1$ such that t_n is a normal form. \square

The usual way to show that a term rewriting system R is terminating is to demonstrate that each rule in R reduces the complexity of terms according to some carefully chosen measure.

Definition 2.6.14 (Confluence). A Σ -term rewriting system R is *confluent* (or is *Church-Rosser*) if whenever $t \rightarrow_R^* t_1$ and $t \rightarrow_R^* t_2$, there is a term t_3 such that $t_1 \rightarrow_R^* t_3$ and $t_2 \rightarrow_R^* t_3$. \square

Definition 2.6.15 (Completeness). A Σ -term rewriting system R is *complete* if it is both terminating and confluent. \square

Completeness of a term rewriting system should not be confused with completeness of a proof system, as in for example Theorem 2.6.10 above.

Exercise 2.6.16. Suppose that R is a complete Σ -term rewriting system, and let $t \in |T_\Sigma(X)|$ be a Σ -term. Show that there is a unique normal form $NF_R(t) \in |T_\Sigma(X)|$ such that $t \rightarrow_R^* NF_R(t)$.

HINT: An *abstract reduction system* consists of a set A together with a binary relation $\rightarrow \subseteq A \times A$. A Σ -term rewriting system R is a particular example, where $A = |T_\Sigma(X)|$ and \rightarrow is \rightarrow_R . Concepts such as normal form and confluence make sense in the context of any abstract reduction system, and the required property holds in this more abstract setting. \square

Example 2.6.17. The term rewriting system RBOOL from Example 2.6.6 is both terminating and confluent, and is therefore complete. As the reduction sequence in Example 2.6.6 shows, $NF_{\text{RBOOL}}(\neg(p \wedge (q \Rightarrow \neg \text{false}))) = \neg p$.

The term rewriting system $\text{RBOOL}' = \text{RBOOL} \cup \{p \wedge q \rightarrow q \wedge p\}$ is not terminating: $p \wedge q \rightarrow_{\text{RBOOL}'} q \wedge p \rightarrow_{\text{RBOOL}'} p \wedge q \rightarrow_{\text{RBOOL}'} q \wedge p \rightarrow_{\text{RBOOL}'} \dots$.

The term rewriting system $\text{RBOOL}'' = \text{RBOOL} \cup \{(p \wedge q) \wedge r \rightarrow p \wedge (q \wedge r)\}$ is not confluent: $(p \wedge \neg p) \wedge q \rightarrow_{\text{RBOOL}''} \text{false} \wedge q$ and $(p \wedge \neg p) \wedge q \rightarrow_{\text{RBOOL}''} p \wedge (\neg p \wedge q)$, and both $\text{false} \wedge q$ and $p \wedge (\neg p \wedge q)$ are normal forms. \square

Exercise 2.6.18. Is your term rewriting system RNAT from Exercise 2.6.7 complete? If not, find an alternative term rewriting system for NAT that is complete. \square

Exercise 2.6.19. A Σ -term rewriting system R is *weakly confluent* if whenever $t \rightarrow_R t_1$ and $t \rightarrow_R t_2$, there is a term t_3 such that $t_1 \rightarrow_R^* t_3$ and $t_2 \rightarrow_R^* t_3$. Find a term rewriting system that is weakly confluent but not confluent. (HINT: Weak confluence plus termination implies confluence, so don't bother looking at terminating term rewriting systems.) Weak confluence is a much easier condition to check than confluence, so the usual way to prove that a term rewriting system is confluent is to show that it is weakly confluent and terminating. \square

In view of the obvious analogy between reduction and computation, $NF_R(t)$ can be thought of as the *value* of t ; since $NF_R(t)$ need not be a ground term, this is a more general notion of computation than the usual one.

Exercise 2.6.20. Convince yourself that $NF_R: |T_\Sigma(X)| \rightarrow |T_\Sigma(X)|$ is computable for any finite complete term rewriting system R — perhaps try to implement it in your favourite programming language. \square

Theorem 2.6.21 (Decision procedure for convertibility). *If R is complete, then $t \sim_R t'$ iff $NF_R(t) = NF_R(t')$.* \square

Exercise 2.6.22. Prove Theorem 2.6.21. HINT: The proof does not depend on the definition of \rightarrow_R , but only on the assumption that R is complete. \square

Since $t \sim_R t'$ iff $\text{Eq}(R) \vdash_\Sigma t = t'$ (by soundness and completeness of convertibility) iff $\text{Eq}(R) \models_\Sigma t = t'$ (by soundness and completeness of the equational calculus), Theorem 2.6.21 constitutes a decision procedure for consequence:

Corollary 2.6.23 (Decision procedure for $\text{Eq}(R) \models_\Sigma t = t'$). *If R is complete, then $\text{Eq}(R) \models_\Sigma t = t'$ iff $NF_R(t) = NF_R(t')$.* \square

Example 2.6.24. Since the term rewriting system RBOOL from Example 2.6.6 is complete (see Example 2.6.17), Corollary 2.6.23 can be used to prove that $\text{Eq}(\text{RBOOL}) \models_{\Sigma_{\text{BOOL}}} \neg(p \wedge (q \Rightarrow \neg\text{false})) = p \Rightarrow (p \wedge \neg p)$: this follows since we have $\text{NF}_{\text{RBOOL}}(\neg(p \wedge (q \Rightarrow \neg\text{false}))) = \neg p = \text{NF}_{\text{RBOOL}}(p \Rightarrow (p \wedge \neg p))$. Since $\text{Cl}_{\Sigma_{\text{BOOL}}}(\mathcal{E}_{\text{BOOL}}) = \text{Cl}_{\Sigma_{\text{BOOL}}}(\text{Eq}(\text{RBOOL}))$, this proves that $\mathcal{E}_{\text{BOOL}} \models_{\Sigma_{\text{BOOL}}} \neg(p \wedge (q \Rightarrow \neg\text{false})) = p \Rightarrow (p \wedge \neg p)$.

Exercise. Give a derivation of $\mathcal{E}_{\text{BOOL}} \vdash_{\Sigma_{\text{BOOL}}} \neg(p \wedge (q \Rightarrow \neg\text{false})) = p \Rightarrow (p \wedge \neg p)$ in the equational calculus. Compare this with the above proof. \square

Exercise 2.6.25. Recall your complete term rewriting system for NAT from Exercise 2.6.18. Relying on Corollary 2.6.23, use this to prove that $\mathcal{E}_{\text{NAT}} \models_{\Sigma_{\text{NAT}}} \text{succ}(\text{succ}(0)) + \text{succ}(n) = \text{succ}(\text{succ}(\text{succ}(n)))$, and that $\mathcal{E}_{\text{NAT}} \not\models_{\Sigma_{\text{NAT}}} \text{succ}(m) + \text{succ}(n) = \text{succ}(\text{succ}(m+n))$. \square

Exercise 2.6.26. Let $t \rightarrow t'$ be a Σ -rewrite rule of sort s . The following restrictions are often imposed:

- $t \notin X_s$; and
- $\text{FV}(t') \subseteq \text{FV}(t)$.

Show that, if these restrictions are imposed on rewrite rules, then Corollary 2.6.23 holds even without the assumption on Σ made at the beginning of this section. (These restrictions seem harmless since almost no complete term rewriting system contains rules that violate them.) \square

Exercise 2.6.27. Equality of terms in the equational theory of a rewriting systems is also decidable under somewhat weaker requirements than those in Corollary 2.6.23. A term rewriting system R is *weakly normalising* if for each term t there is a finite reduction sequence in R leading from t to a normal form. R is *semi-complete* if it is weakly normalising and confluent.

Generalising Exercise 2.6.16, show that if R is a semi-complete Σ -term rewriting system, then for any Σ -term $t \in |T_{\Sigma}(X)|$ there is a unique normal form $\text{NF}_R(t) \in |T_{\Sigma}(X)|$ such that $t \rightarrow_R^* \text{NF}_R(t)$. Moreover, convince yourself that the function $\text{NF}_R: |T_{\Sigma}(X)| \rightarrow |T_{\Sigma}(X)|$ is then computable. Finally, show that the property captured by Corollary 2.6.23 holds for all semi-complete term rewriting systems R . \square

By Corollary 2.6.23, the problem of deciding consequence $\mathcal{E} \models_{\Sigma} e$ is reduced to the problem of finding a finite complete term rewriting system R such that $\text{Cl}_{\Sigma}(\mathcal{E}) = \text{Cl}_{\Sigma}(\text{Eq}(R))$. Clearly, by Theorem 2.4.15, this is not always possible. But the *Knuth-Bendix completion algorithm* can sometimes be used to produce such an R given \mathcal{E} together with an order relation on terms. The algorithm works by pinpointing causes of failure of (weak) confluence and adding rules to correct them, where the supplied term ordering is used to orient these new rules. The algorithm is iterative and may fail to terminate; it may also fail because the ordering supplied is inadequate.

The Knuth-Bendix completion algorithm can also be used to reason about initial models of specifications, using a method known as *inductionless induction* or *proof by consistency*. This method is based on the observation that an equation $t = t'$

holds in the initial models of $\langle \Sigma, \mathcal{E} \rangle$ iff there is no ground equation $u = u'$ such that $\mathcal{E} \not\models u = u'$ and $\mathcal{E} \cup \{t = t'\} \models u = u'$. (**Exercise:** Prove this fact.) Given a complete term rewriting system R such that $Cl_{\Sigma}(\mathcal{E}) = Cl_{\Sigma}(Eq(R))$ (perhaps produced using the Knuth-Bendix algorithm), the Knuth-Bendix algorithm is used to produce a complete term rewriting system R' for $\mathcal{E} \cup \{t = t'\}$ by extending R . It is then possible to test if R and R' have the same normal forms for ground Σ -terms; if so, then $t = t'$ holds in the initial models of $\langle \Sigma, \mathcal{E} \rangle$.

2.7 Fiddling with the definitions

In principle, the specification framework presented in the preceding sections is powerful enough for any conceivable computational application. This is made precise by a theorem in [BT87] (cf. [Vra88]) which states that for every reachable *semi-computable* Σ -algebra A there is a presentation $\langle \Sigma', \mathcal{E}' \rangle$ with finite \mathcal{E}' such that $A = A' |_{\Sigma}$ for some initial model $A' \in IMod[\langle \Sigma', \mathcal{E}' \rangle]$. (See [BT87] for the definition of semi-computable algebra.) In spite of this fact, there are several reasons why this framework is inconvenient for use in practice.

One deficiency becomes apparent as soon as one attempts to write specifications that are somewhat larger than the examples we have seen so far. In order to be understandable and usable, large specifications must be built up incrementally from smaller specifications. Specification mechanisms designed to cope with such problems of scale are presented in Chapter 5. These methods also solve the problem illustrated by Exercise 2.5.20; see Exercise 5.1.11.

Another difficulty arises from the relatively low level of equational logic as a language for describing constraints to be satisfied by the operations of an algebra. When using equational axioms, it is often necessary to write a dozen equations to express a property that can be formulated much more clearly using a single axiom in some more powerful logic. Some properties that are easy to express in more powerful systems are not expressible at all using equations. Similar awkwardness is caused by the limitations of the type system used here, in comparison with the polymorphic type systems of modern programming languages such as Standard ML [Pau96]. Finally, the present framework is only able to cope conveniently with algebras comprised of *total* and *deterministic* functions operating on data values built by *finitary* compositions of such functions, a limitation which rules out its use for very many programs of interest.

All these difficulties can be addressed by making appropriate modifications to the standard framework presented in the preceding sections. An example was already given in Section 1.5.2 where it was shown how signature morphisms could be replaced by derived signature morphisms. This section is devoted to a sketch of some other possible modifications. The presentation is very brief and makes no attempt to be truly comprehensive; the interested reader will find further details (and further citations) in the cited references.

2.7.1 Conditional equations

The most obvious kind of modification is to replace the use of equational axioms by formulae in a more expressive language. Some care is required since a number of the results presented above depend on the use of equational axioms. A relatively unproblematic choice is to use equations that apply only when certain pre-conditions (expressed as equations) are satisfied.

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

Definition 2.7.1 (Conditional equation). A *conditional Σ -equation* $\forall X \bullet t_1 = t'_1 \wedge \dots \wedge t_n = t'_n \Rightarrow t_0 = t'_0$ consists of

- a finite S -sorted set X (of variables), such that $X_s \subseteq \mathcal{X}$ for all $s \in S$; and
- for each $0 \leq j \leq n$ (where $n \geq 0$), two Σ -terms $t_j, t'_j \in |T_\Sigma(X)|_{s_j}$ for some sort $s_j \in S$.

A Σ -algebra A *satisfies* a conditional Σ -equation $\forall X \bullet t_1 = t'_1 \wedge \dots \wedge t_n = t'_n \Rightarrow t_0 = t'_0$ if for every (S -sorted) function $v: X \rightarrow |A|$, if $(t_1)_A(v) = (t'_1)_A(v)$ and \dots and $(t_n)_A(v) = (t'_n)_A(v)$, then $(t_0)_A(v) = (t'_0)_A(v)$. \square

Note that variables in the conditions $(t_1 = t'_1 \wedge \dots \wedge t_n = t'_n)$ that do not appear in the consequent $(t_0 = t'_0)$ can be seen as existentially quantified: for example, the conditional equation $\forall a, b: t \bullet a \times b = 1 \Rightarrow a \times a^{-1} = 1$ is equivalent to the formula $\forall a: t \bullet (\exists b: t \bullet a \times b = 1) \Rightarrow a \times a^{-1} = 1$ in ordinary first-order logic.

Exercise 2.7.2. Define the translation of conditional Σ -equations by a signature morphism $\sigma: \Sigma \rightarrow \Sigma'$. \square

The remaining definitions of Sections 2.1–2.5 require only superficial changes, and most results go through with appropriate modifications.

Let $\langle \Sigma, \mathcal{E} \rangle$ be a presentation, where \mathcal{E} is a set of conditional Σ -equations. $\text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$ is not always a variety, as is (almost) shown by Example 2.2.11; in this sense, the power of conditional equations is strictly greater than that of ordinary equations.

Exercise 2.7.3. The cancellation law given in Example 2.2.11 is not a conditional equation. Give a version of this example that uses only conditional equations. HINT: Equality can be axiomatised as an operation $eq: s \times s \rightarrow \text{Bool}$. \square

In spite of this increase in expressive power, there is a proof system that is sound and complete with respect to conditional equational consequence [Sel72], and the quotient construction can be used to construct an initial model of $\langle \Sigma, \mathcal{E} \rangle$ [MT92] (cf. Lemma 3.3.12 below). Term rewriting with conditional rewrite rules is possible, but there are some complications; see [Klo92] and [Mid93].

Exercise 2.7.4. [Sel72] gives a proof system that is sound and complete for conditional equational consequence in the single-sorted case. Extend this to the many-sorted case, where explicit quantifiers are required for the same reason as in the equational calculus. \square

Exercise 2.7.5. Recall Exercise 2.5.21 concerning the specification of a function $ch: Nat \rightarrow Nat$ that for each natural number n chooses an arbitrary number that is greater than n . Modify this, using a conditional equation to make ch choose an arbitrary number that is *less* than n when $0 < n$. \square

Example 2.7.6. Let $HA = \langle \Sigma_{HA}, \mathcal{E}_{HA} \rangle$ be the following presentation.³

spec $HA = \text{sorts } Bool$

ops $\text{true}: Bool$

$\text{false}: Bool$

$\neg _ : Bool \rightarrow Bool$

$_ \vee _ : Bool \times Bool \rightarrow Bool$

$_ \wedge _ : Bool \times Bool \rightarrow Bool$

$_ \Rightarrow _ : Bool \times Bool \rightarrow Bool$

$\forall p, q, r: Bool$

• $p \vee (q \vee r) = (p \vee q) \vee r$

• $p \wedge (q \wedge r) = (p \wedge q) \wedge r$

• $p \vee q = q \vee p$

• $p \wedge q = q \wedge p$

• $p \vee (p \wedge q) = p$

• $p \wedge (p \vee q) = p$

• $p \vee \text{true} = \text{true}$

• $p \vee \text{false} = p$

• $(p \vee (r \wedge q) = p) \Rightarrow ((q \Rightarrow p) \vee r = (q \Rightarrow p))$

• $((q \Rightarrow p) \vee r = (q \Rightarrow p)) \Rightarrow (p \vee (r \wedge q) = p)$

• $\neg p = (p \Rightarrow \text{false})$

Models of HA are called *Heyting algebras*.

Exercise. Recall the presentation BA of Boolean algebras in Example 2.2.4. Show that every Boolean algebra is a Heyting algebra. Then repeat the exercise in Example 2.2.4, building for every Heyting algebra H a lattice $\langle |H|, \leq_H \rangle$ with top and bottom elements. Check that the conditional axioms concerning the implication \Rightarrow can now be captured by requiring that $r \wedge q \leq_H p$ is equivalent to $r \leq_H q \Rightarrow p$. Show that the lattice is distributive.

Give an example of a Heyting algebra that is not Boolean. Check which of the axioms of the presentation BA do not follow from HA.

Prove that an *equational* presentation with the same models as HA can be given. **HINT:** Use Theorem 2.2.10. Or consider the following properties of the implication: $p \Rightarrow p = \text{true}$, $q \wedge (q \Rightarrow p) = q \wedge p$, $p \vee (q \Rightarrow p) = q \Rightarrow p$, and $q \Rightarrow (p \wedge r) = (q \Rightarrow p) \wedge (q \Rightarrow r)$. \square

³ We use the same symbol \Rightarrow for implication in conditional equations and for an operation in the presentation below — the usual symbols are used for other propositional connectives as well, as in Example 2.2.4. We use extra space around the implication symbol in the conditional equations below in order to make them easier to read.

2.7.2 Reachable semantics

In Section 2.5, the motivation given for taking a presentation $\langle \Sigma, \mathcal{E} \rangle$ to denote the class $IMod[\langle \Sigma, \mathcal{E} \rangle]$ of its initial models was the desire to exclude models containing junk and confusion. The need to exclude models containing confusion stems mainly from the use of equational axioms, which make it impossible to rule out degenerate models having a single value of each sort in Σ . If a more expressive language is used for axioms, or if degenerate models are ruled out by some other means, then models containing confusion need not be excluded.

Example 2.7.7. Consider the following specification of sets of natural numbers (a variant of the one in Exercise 2.5.23):

```

spec SETNAT = sorts Bool, Nat, NatSet
                ops true: Bool
                    false: Bool
                    -- ∨ --: Bool × Bool → Bool
                    0: Nat
                    succ: Nat → Nat
                    eq: Nat × Nat → Bool
                    ∅: NatSet
                    add: Nat × NatSet → NatSet
                    -- ∈ --: Nat × NatSet → Bool
                ∀ p: Bool, m, n: Nat, S: NatSet
                • p ∨ true = true
                • p ∨ false = p
                • eq(n, n) = true
                • eq(0, succ(n)) = false
                • eq(succ(n), 0) = false
                • eq(succ(m), succ(n)) = eq(m, n)
                • n ∈ ∅ = false
                • m ∈ add(n, S) = eq(m, n) ∨ m ∈ S

```

There are many different models of SETNAT, including algebras having a single value of each sort. Suppose we restrict attention to algebras that do not satisfy the equation $true = false$; this excludes such degenerate models (see the exercise below). Consider the following two equations:

Commutativity of *add*:

$$\forall m, n: Nat, S: NatSet \bullet add(m, add(n, S)) = add(n, add(m, S))$$

Idempotency of *add*:

$$\forall n: Nat, S: NatSet \bullet add(n, add(n, S)) = add(n, S)$$

The models of SETNAT that do not satisfy $true = false$ may be classified according to which of these two equations they satisfy:

“List-like” algebras: *add* is neither commutative nor idempotent.

“Set-like” algebras: *add* is both commutative and idempotent.

“Multiset-like” algebras: *add* is commutative but not idempotent.

“List-like” algebras without repeated adjacent entries: *add* is idempotent but not commutative.

There are also “hybrid” models of SETNAT , e.g. those in which *add* is commutative but is only idempotent for $n \neq 0$. The initial models of SETNAT are “list-like” algebras. Adding the commutativity and idempotency requirements to SETNAT as additional axioms would eliminate all but the “set-like” algebras.

Exercise. Show that restricting attention to models of SETNAT that do not satisfy $\text{true} = \text{false}$ eliminates all but “sensible” realisations of sets of natural numbers, by forcing $\text{eq}(\text{succ}^m(0), \text{succ}^n(0)) = \text{true}$ iff $m = n$ iff $\text{succ}^m(0) = \text{succ}^n(0)$, and $a \in \text{add}(a_1, \text{add}(a_2, \dots, \text{add}(a_p, \emptyset) \dots)) = \text{true}$ iff $\text{eq}(a, a_1) = \text{true}$ or \dots or $\text{eq}(a, a_p) = \text{true}$, for $m, n, p \geq 0$. Note that m, n and p are ordinary integers here, *not* values of the sort *Nat*, and $\text{succ}^m(0)$ means $\underbrace{\text{succ}(\dots \text{succ}(0) \dots)}_{m \text{ times}}$. \square

Consideration of examples like the one above suggests various alternatives to taking the initial semantics of specifications. One choice is to require signatures to include the sort *Bool* and the constants *true* and *false*, and to exclude models satisfying $\text{true} = \text{false}$. This might be called taking the *non-degenerate loose semantics* of specifications. Another choice is to additionally exclude models containing junk:

Definition 2.7.8 (Reachable semantics). Let $\Sigma = \langle S, \Omega \rangle$ be a signature such that *Bool* $\in S$ and $\text{true}:\text{Bool}$ and $\text{false}:\text{Bool}$ are in Ω . A *reachable non-degenerate model* of a presentation $\langle \Sigma, \mathcal{E} \rangle$ is a reachable Σ -algebra A such that $A \models_{\Sigma} \mathcal{E}$ and $A \not\models_{\Sigma} \text{true} = \text{false}$. $\text{RMod}[\langle \Sigma, \mathcal{E} \rangle]$ is the class of all reachable non-degenerate models of $\langle \Sigma, \mathcal{E} \rangle$. Taking $\langle \Sigma, \mathcal{E} \rangle$ to denote $\text{RMod}[\langle \Sigma, \mathcal{E} \rangle]$ is called taking its *reachable semantics*. \square

The motivation for excluding models containing junk is the same as in the case of initial semantics. $\text{RMod}[\langle \Sigma, \mathcal{E} \rangle]$ is not always an isomorphism class of models, as Example 2.7.7 demonstrates (the classification given there was for *all* models that do not satisfy $\text{true} = \text{false}$, but it also applies to the reachable models in this class). There is still a problem when operations are not defined in a sufficiently complete way, although the problem is less severe than in the case of initial semantics.

Exercise 2.7.9. Reconsider the problem posed in Exercise 2.5.20, by writing a reachable model specification of natural numbers including a subtraction operation $-- -: \text{Nat} \times \text{Nat} \rightarrow \text{Nat}$ together with the axioms $\forall m:\text{Nat} \bullet m - 0 = m$ and $\forall m, n:\text{Nat} \bullet \text{succ}(m) - \text{succ}(n) = m - n$. Recall from Exercise 2.5.20 the assumption that we are willing to accept any value for $m - n$ when $n > m$, which is why the axioms do not constrain the value of $m - n$ in this case. List some of the reachable non-degenerate models of this specification, and decide whether the models you considered in Exercise 2.5.20 are reachable non-degenerate models (ignoring the difference in signatures). From an intuitive point of view, is this an adequate class of models for this specification? \square

Exercise 2.7.10. Definition 2.7.8 permits algebras $A \in RMod[\langle \Sigma, \mathcal{E} \rangle]$ with values of sort *Bool* other than $true_A$ and $false_A$. This is ruled out if all operations delivering results in sort *Bool* are defined in a sufficiently complete way to yield either *true* or *false* on each argument that is definable by a ground term. Check that the specification SETNAT in Example 2.7.7 ensures this property, and so all of its reachable non-degenerate models have a two-element carrier of sort *Bool*. Give an example of a specification for which this is not the case. \square

The equational calculus is sound for reasoning about the reachable semantics of presentations, since $RMod[\langle \Sigma, \mathcal{E} \rangle] \subseteq Mod[\langle \Sigma, \mathcal{E} \rangle]$ for any presentation $\langle \Sigma, \mathcal{E} \rangle$. It is sound to add induction rule schemes such as those given in Section 2.5; these are sound for any class of reachable models. Completeness is unachievable, for exactly the same reason as in the case of initial semantics; the proof of Theorem 2.5.26 can be repeated in this context almost without change. Finally, the techniques of term rewriting presented in Section 2.6 remain sound.

Initial semantics cannot be used for specifications with axioms that are more expressive than (infinitary) conditional equations [Tar86b], in the sense that initial models of such specifications are not guaranteed to exist. To illustrate the problem, the following example shows what can go wrong when the language of axioms is extended to permit disjunctions of equations.

Example 2.7.11. Consider the following specification:

spec STATUS = **sorts** *Status*
ops *single*:*Status*
married:*Status*
widowed:*Status*
• $widowed = single \vee widowed = married$

where disjunction of equations has the obvious interpretation. There are three kinds of algebras in $Mod[STATUS]$:

1. Those satisfying $single = widowed = married$.
2. Those satisfying $single = widowed \neq married$.
3. Those satisfying $single \neq widowed = married$.

None of these is an initial model of STATUS: there are no homomorphisms from algebras in the first class to algebras in either of the other two classes, and no homomorphisms in either direction between algebras in the second and third classes. \square

In contrast, reachable semantics can be used for specifications with axioms of any form (once a definition of satisfaction of such axioms by algebras has been given, of course).

Another alternative to initial semantics deserves brief mention.

Definition 2.7.12 (Final semantics). Let $\Sigma = \langle S, \Omega \rangle$ be a signature such that *Bool* $\in S$ and $true:Bool$ and $false:Bool$ are in Ω . A Σ -algebra $A \in RMod[\langle \Sigma, \mathcal{E} \rangle]$ is a *final* (or *terminal*) *model* of $\langle \Sigma, \mathcal{E} \rangle$ if for every $B \in RMod[\langle \Sigma, \mathcal{E} \rangle]$ there is a unique Σ -homomorphism $h: B \rightarrow A$. Taking $\langle \Sigma, \mathcal{E} \rangle$ to denote the class of its final models is called taking its *final semantics*. \square

As in the case of initial semantics, the final models of a presentation form an isomorphism class. Recall that a model of a presentation is initial iff it contains no junk and no confusion (Exercise 2.5.19). We can give a similar characterisation of final models as the models containing no junk and *maximal confusion*: a final model A satisfies as many ground equations as possible, subject to the restriction that $A \not\models \text{true} = \text{false}$ (imposed on all reachable non-degenerate models).

Example 2.7.13. Recall the specification SETNAT from Example 2.7.7, and the classification of models of SETNAT according to the commutativity and idempotence of *add*. The final models of SETNAT are in the class of “set-like” algebras, in which *add* is both commutative and idempotent. (**Exercise:** Why?) \square

Not all presentations with equational axioms have final models, but it is possible to impose conditions on the form of presentations that guarantee the existence of final models [BDP⁺79].

Exercise 2.7.14. Find a variation on the specification STATUS in Example 2.7.11 that has no final models. \square

When reachable or final semantics of presentations is used with equational or conditional equational axioms, sometimes more operations are required in specifications than in the case of initial semantics. These additional operations are needed to provide ways of “observing” values of sorts other than *Bool*, in order to avoid models that are degenerate on these other sorts. For example, the presence of the operation *eq* in Example 2.7.7 ensures that $\text{succ}^m(0) = \text{succ}^n(0)$ only if $m = n$ in all models that do not satisfy $\text{true} = \text{false}$; it would not be needed if we were interested only in the initial models of SETNAT . Such operations are not required if inequations are allowed as axioms.

Exercise 2.7.15. Recall the presentation NAT given in Exercise 2.5.4. Augment this with the sort *Bool* and constants *true, false: Bool* (to make reachable and final semantics applicable), and show that final models of the resulting specification have a single value of sort *Nat*. Add an operation $\text{even}: \text{Nat} \rightarrow \text{Bool}$, with the following axioms:

$\forall n: \text{Nat}$

- $\text{even}(\text{succ}(\text{succ}(n))) = \text{even}(n)$
- $\text{even}(\text{succ}(0)) = \text{false}$
- $\text{even}(0) = \text{true}$

Show that final models of the resulting specification have exactly two values of sort *Nat*. Replace *even* with $\leq: \text{Nat} \times \text{Nat} \rightarrow \text{Bool}$, with appropriate axioms, and show that final models of the resulting specification satisfy $\text{succ}^m(0) = \text{succ}^n(0)$ iff $m = n$. (We have already seen that this is the case if $\text{eq}: \text{Nat} \times \text{Nat} \rightarrow \text{Bool}$ is added in place of \leq .) \square

Although the inclusion of additional operations tends to make specifications longer, it is not an artificial device. In practice, one would expect each sort to come with an assortment of operations for creating, manipulating and observing values of that sort, so specifications such as NAT are less natural than NAT augmented with operations like \leq and/or *eq*.

2.7.3 Dealing with partial functions: error algebras

An obvious inadequacy of the framework(s) presented above stems from the use of *total* functions in algebras to interpret the operation names in a signature. Since partial functions are not at all uncommon in computer science applications — a very simple example being the predecessor function $pred: Nat \rightarrow Nat$, which is undefined on 0 — a great deal of work has gone into ways of lifting this restriction. Three main approaches are discussed below:

Error algebras (this subsection): Predecessor is regarded as a total function, with $pred(0)$ specified to yield an *error* value.

Partial algebras (Section 2.7.4): Predecessor is regarded as a partial function.

Order-sorted algebras (Section 2.7.5): Predecessor is regarded as a total function on a sub-domain that excludes the value 0.

A fourth approach is to use ordinary (total) algebras, leaving the value of $pred(0)$ unspecified. This is more an attempt to avoid the issue than a solution, and it is workable only in frameworks that deal adequately with definitions that are not sufficiently complete; see Exercises 2.5.20, 2.7.9, and 5.1.11.

The most obvious way of adding error values to algebras does not work, as the following example demonstrates.

Example 2.7.16. Consider the following specification of the natural numbers, where $pred(0)$ is specified to yield an error:

```
spec NATPRED = sorts Nat
ops  0: Nat
      succ: Nat → Nat
      pred: Nat → Nat
      error: Nat
      ++: Nat × Nat → Nat
      -- × --: Nat × Nat → Nat
∀m, n: Nat
  • pred(succ(n)) = n
  • pred(0) = error
  • 0 + n = n
  • succ(m) + n = succ(m + n)
  • 0 × n = 0
  • succ(m) × n = (m × n) + n
```

Initial models of NATPRED will have many “non-standard” values of sort *Nat*, in addition to the intended one (*error*). For example, the axioms of NATPRED do not force the ground terms $pred(error)$ and $pred(error) + 0$ to be equal to any “normal” value, or to *error*. (**Exercise:** Give an initial model of NATPRED.) A possible solution to this is to add axioms that collapse these non-standard values to a single point:

```

spec NATPRED = sorts Nat
                ops ...
                 $\forall m, n: \text{Nat}$ 
                • ...
                •  $\text{succ}(\text{error}) = \text{error}$ 
                •  $\text{pred}(\text{error}) = \text{error}$ 
                •  $\text{error} + n = \text{error}$ 
                •  $n + \text{error} = \text{error}$ 
                •  $\text{error} \times n = \text{error}$ 
                •  $n \times \text{error} = \text{error}$ 

```

Unfortunately, NATPRED now has only trivial models: $\text{error} = 0 \times \text{error} = 0$, and so $\text{error} = \text{succ}(\text{error}) = \text{succ}(0)$, $\text{error} = \text{succ}(\text{error}) = \text{succ}(\text{succ}(0))$, and so on. \square

The above example suggests that a more delicate treatment is required. A number of approaches have been proposed; here we follow [GDLE84], which is fairly powerful without sacrificing simplicity and elegance. The main ideas of this approach are:

- Error values are distinguished from non-error (“OK”) values.
- In an *error signature*, operations that may produce errors when given OK arguments (*unsafe* operations) are distinguished from those that always preserve OK-ness (*safe* operations).
- In an *error algebra*, each carrier is partitioned into an error part and an OK part. Safe operations are required to produce OK results for OK arguments, and homomorphisms are required to preserve OK-ness.
- In equations, variables that can take OK values only (*safe* variables) are distinguished from variables that can take any value (*unsafe* variables). Assignments of values to variables are required to map safe variables to OK values.

Definition 2.7.17 (Error signature). An *error signature* is a triple $\Sigma = \langle S, \Omega, \text{safe} \rangle$ where

- $\langle S, \Omega \rangle$ is an ordinary signature; and
- *safe* is an $S^* \times S$ -sorted set of functions $\langle \text{safe}_{w,s}: \Omega_{w,s} \rightarrow \{tt, ff\} \rangle_{w \in S^*, s \in S}$.

An operation $f: s_1 \times \dots \times s_n \rightarrow s$ in Σ is *safe* if $\text{safe}_{s_1 \dots s_n, s}(f) = tt$; otherwise it is *unsafe*. \square

Example 2.7.16 (revisited). An appropriate error signature for NATPRED would be the following:

```

 $\Sigma$ NATPRED = sorts Nat
                ops 0: Nat
                 $\text{succ}: \text{Nat} \rightarrow \text{Nat}$ 
                 $\text{pred}: \text{Nat} \rightarrow \text{Nat}, \text{unsafe}$ 
                 $\text{error}: \text{Nat}, \text{unsafe}$ 
                 $-- + --: \text{Nat} \times \text{Nat} \rightarrow \text{Nat}$ 
                 $-- \times --: \text{Nat} \times \text{Nat} \rightarrow \text{Nat}$ 

```

Obviously, *error* is unsafe, and *pred* is unsafe since it produces an error when applied to 0; all the remaining operations are safe. (By convention, the safe operations are those that are not explicitly marked as unsafe.) \square

In the rest of this section, let $\Sigma = \langle S, \Omega, \text{safe} \rangle$ be an error signature.

Definition 2.7.18 (Error algebra). An error Σ -algebra A consists of

- an ordinary Σ -algebra A ; and
- an S -sorted set of functions $OK = \langle OK_s: |A|_s \rightarrow \{tt, ff\} \rangle_{s \in S}$

such that safe operations preserve OK-ness: for every $f: s_1 \times \dots \times s_n \rightarrow s$ in Σ such that $\text{safe}_{s_1 \dots s_n, s}(f) = tt$ and $a_1 \in |A|_{s_1}, \dots, a_n \in |A|_{s_n}$ such that $OK_{s_1}(a_1) = \dots = OK_{s_n}(a_n) = tt$, $OK_s(f_A(a_1, \dots, a_n)) = tt$. A value $a \in |A|_s$ for $s \in S$ is an *OK value* if $OK_s(a) = tt$; otherwise it is an *error value*. \square

Definition 2.7.19 (Error homomorphism). Let A and B be error Σ -algebras. An error Σ -homomorphism $h: A \rightarrow B$ is an S -sorted function $h: |A| \rightarrow |B|$ with the usual homomorphism property (for all $f: s_1 \times \dots \times s_n \rightarrow s$ in Σ and $a_1 \in |A|_{s_1}, \dots, a_n \in |A|_{s_n}$, $h_s(f_A(a_1, \dots, a_n)) = f_B(h_{s_1}(a_1), \dots, h_{s_n}(a_n))$) such that h preserves OK-ness: for every $s \in S$ and $a \in |A|_s$ such that $OK_s(a) = tt$ (in A), $OK_s(h_s(a)) = tt$ (in B). \square

Definition 2.7.20 (Error variable set). An error S -sorted variable set X consists of an S -sorted set X such that $X_s \subseteq \mathcal{X}$ for all $s \in S$, and an S -sorted set of functions $\text{safe} = \langle \text{safe}_s: X_s \rightarrow \{tt, ff\} \rangle_{s \in S}$. A variable $x: s$ in X is *safe* if $\text{safe}_s(x) = tt$; otherwise it is *unsafe*. An *assignment* of values in an error Σ -algebra A to an error S -sorted variable set X is an S -sorted function $v: X \rightarrow |A|$ assigning OK values to safe variables: for every $x: s$ in X such that $\text{safe}_s(x) = tt$, $OK_s(v_s(x)) = tt$. \square

Definition 2.7.21 (Error algebra of terms). Let X be an error S -sorted variable set. The error Σ -algebra $ET_\Sigma(X)$ of terms with variables X is defined in an analogous way to the ordinary term algebra $T_\Sigma(X)$, with the following partition of the S -sorted set of terms into OK and error values:

For all sorts $s \in S$ and Σ -terms $t \in |ET_\Sigma(X)|_s$, if t contains an unsafe variable or operation then $OK_s(t) = ff$; otherwise $OK_s(t) = tt$.

We adopt the same notational conventions for terms as before, dropping sort decorations, etc., when there is no danger of confusion. Let ET_Σ denote $ET_\Sigma(\emptyset)$. \square

The definitions of term evaluation, error equation, satisfaction of an error equation by an error algebra, error presentation, model of an error presentation, semantic consequence, and initial model are analogous to the definitions given earlier in the standard many-sorted algebraic framework (Definitions 1.4.5, 2.1.1, 2.1.2, 2.2.1, 2.2.2, 2.3.6 and 2.5.13 respectively). Because assignments are required to map safe variables to OK values, an error equation may be satisfied by an error algebra even if it is not satisfied when error values are substituted for safe variables.

Exercise 2.7.22. Spell out the details of these definitions. \square

As before, every error presentation has an isomorphism class of initial models, and an analogous quotient construction gives an initial model.

Definition 2.7.23 (Congruence generated by a set of equations). Let \mathcal{E} be a set of error Σ -equations. The Σ -congruence $\equiv_{\mathcal{E}}$ on ET_{Σ} is defined by $t \equiv_{\mathcal{E}} t' \iff \mathcal{E} \models_{\Sigma} t = t'$ for all $t, t' \in |ET_{\Sigma}|$. $\equiv_{\mathcal{E}}$ is called the Σ -congruence generated by \mathcal{E} . (NOTE: A Σ -congruence on an error Σ -algebra A is just an ordinary Σ -congruence on the ordinary Σ -algebra underlying A .) \square

Definition 2.7.24 (Quotient error algebra). Let A be an error Σ -algebra, and let \equiv be a Σ -congruence on A . The definition of A/\equiv , the *quotient error algebra of A modulo \equiv* , is analogous to that of the ordinary quotient algebra A/\equiv , with the following partition of congruence classes into OK and error values:

For all sorts $s \in S$ and congruence classes $[a]_{\equiv_s} \in |A/\equiv|_s$, if there is some $b \in [a]_{\equiv_s}$ such that $OK_s(b) = tt$ (in A), then $OK_s([a]_{\equiv_s}) = tt$ (in A/\equiv); otherwise $OK_s([a]_{\equiv_s}) = ff$. \square

Note that if there are both OK and error values in a congruence class, the class is regarded as an OK value in the quotient.

Theorem 2.7.25 (Initial model theorem). *The error Σ -algebra $ET_{\Sigma}/\equiv_{\mathcal{E}}$ is an initial model of the error presentation $\langle \Sigma, \mathcal{E} \rangle$.* \square

Exercise 2.7.26. Sketch a proof of Theorem 2.7.25. HINT: Take inspiration from the proof of Theorem 2.5.14. \square

Exercise 2.7.27. Try to find conditions analogous to “no junk” and “no confusion” that characterise the initial models of an error presentation. \square

Example 2.7.16 (revisited). Using the approach outlined above, here is an improved version of the specification NATPRED:

```

spec NATPRED = sorts Nat
                ops  0: Nat
                    succ: Nat  $\rightarrow$  Nat
                    pred: Nat  $\rightarrow$  Nat, unsafe
                    error: Nat, unsafe
                    ++: Nat  $\times$  Nat  $\rightarrow$  Nat
                    --  $\times$  -: Nat  $\times$  Nat  $\rightarrow$  Nat
                 $\forall m, n: \text{Nat}$ 
                    • pred(succ(n)) = n
                    • pred(0) = error
                    • 0 + n = n
                    • succ(m) + n = succ(m + n)
                    • 0  $\times$  n = 0
                    • succ(m)  $\times$  n = (m  $\times$  n) + n

```

(By convention, variables in equations are safe unless otherwise indicated.) In initial models of NATPRED, the error values of sort *Nat* correspond exactly to “error messages”, i.e. ground terms containing at least one occurrence of *error*. These terms can be regarded as recording the sequence of events that took place since the error occurred. The record is accurate since the initial models of NATPRED do *not* satisfy equations like $0 \times error = 0$, in contrast to the initial models of the earlier version. To collapse the error values to a single point without affecting the OK values, axioms can be added as follows:

spec NATPRED = **sorts** *Nat*
ops ...
 $\forall m, n, k: Nat, k: Nat: unsafe$
 • ...
 • $pred(error) = error$
 • $succ(error) = error$
 • $error + k = error$
 • $k + error = error$
 • $error \times k = error$
 • $k \times error = error$

It is also possible to specify *error recovery* using this approach:

spec NATPRED = **sorts** *Nat*
ops ...
 $recover: Nat \rightarrow Nat$
 $\forall m, n, k: Nat, k: Nat: unsafe$
 • ...
 • $recover(error) = 0$
 • $recover(n) = n$

In initial models of this version of NATPRED, *recover* is the identity on *Nat* except that $recover(error)$ gives the OK value 0. \square

Although only initial semantics of error presentations has been mentioned above, the alternatives of reachable and final semantics apply as in the standard case. The key points of the standard framework not mentioned here (e.g. analogues to the soundness, completeness and incompleteness theorems) carry over to the present framework as well.

Exercise 2.7.28. Find a definition of error signature morphism which makes the Satisfaction Lemma hold, taking the natural definition of the σ -reduct $A'|_{\sigma}$ of an error Σ' -algebra A' induced by an error signature morphism $\sigma: \Sigma \rightarrow \Sigma'$. \square

Although the approach to error specification presented above is quite attractive, there are examples that cannot be treated in this framework.

Exercise 2.7.29. Consider the following specification of *bounded natural numbers*:

spec BOUNDEDNAT = **sorts** *Nat*
ops $0: \text{Nat}$
 $\text{succ}: \text{Nat} \rightarrow \text{Nat}, \text{unsafe}$
 $\text{overflow}: \text{Nat}, \text{unsafe}$
• $\text{succ}(\text{succ}(\text{succ}(\text{succ}(\text{succ}(0)))))) = \text{overflow}$

The intention is to specify a (very) restricted subset of the natural numbers, where an attempt to compute a number larger than 5 results in overflow. Show that an initial model of BOUNDEDNAT will have only one OK value. Change BOUNDEDNAT so that its initial models have six OK values (corresponding to $0, \text{succ}(0), \dots, \text{succ}^5(0)$). What if the bound is 2^{32} rather than 5? \square

2.7.4 Dealing with partial functions: partial algebras

An obvious way to deal with partial functions is to simply change the definition of algebra to allow operation names to be interpreted as partial functions. But for many of the basic notions in the framework that depend on the definition of algebra, beginning with the concepts of subalgebra and homomorphism, there are several ways to extend the usual definition to the partial case. Choosing a coherent combination of these definitions is a delicate matter. Here we follow the approach of [BW82b].

Throughout this section, let $\Sigma = \langle S, \Omega \rangle$ be a signature.

Definition 2.7.30 (Partial algebra). A *partial* Σ -algebra A is like an ordinary Σ -algebra, except that each $f: s_1 \times \dots \times s_n \rightarrow s$ in Σ is interpreted as a *partial* function $(f: s_1 \times \dots \times s_n \rightarrow s)_A: |A|_{s_1} \times \dots \times |A|_{s_n} \rightarrow |A|_s$. The (total) Σ -algebra underlying A is the Σ -algebra A_\perp defined as follows:

- $|A_\perp|_s = |A|_s \uplus \{\perp_s\}$ for every $s \in S$; and
 - $(f: s_1 \times \dots \times s_n \rightarrow s)_{A_\perp}(a_1, \dots, a_n) = \begin{cases} \perp_s & \text{if } a_j = \perp_{s_j} \text{ for some } 1 \leq j \leq n \\ (f: s_1 \times \dots \times s_n \rightarrow s)_A(a_1, \dots, a_n) & \text{if this is defined} \\ \perp_s & \text{otherwise} \end{cases}$
- for every $f: s_1 \times \dots \times s_n \rightarrow s$ and $a_1 \in |A_\perp|_{s_1}, \dots, a_n \in |A_\perp|_{s_n}$. \square

We employ the same notational conventions as before. Note that according to this definition, the value of a constant need not be defined: a constant $c: s$ is associated in an algebra A with a partial function $c_A: \{\langle \rangle\} \rightarrow |A|_s$, where $\{\langle \rangle\}$ is the nullary Cartesian product.

Definition 2.7.31 (Homomorphism). Let A and B be partial Σ -algebras. A *weak* Σ -homomorphism $h: A \rightarrow B$ is an S -sorted (total) function $h: |A| \rightarrow |B|$ such that for all $f: s_1 \times \dots \times s_n \rightarrow s$ in Σ and $a_1 \in |A|_{s_1}, \dots, a_n \in |A|_{s_n}$,

$$\text{if } f_A(a_1, \dots, a_n) \text{ is defined then } f_B(h_{s_1}(a_1), \dots, h_{s_n}(a_n)) \text{ is defined, and} \\ h_s(f_A(a_1, \dots, a_n)) = f_B(h_{s_1}(a_1), \dots, h_{s_n}(a_n)).$$

If moreover h satisfies the condition

if $f_B(h_{s_1}(a_1), \dots, h_{s_n}(a_n))$ is defined then $f_A(a_1, \dots, a_n)$ is defined

then h is called a *strong Σ -homomorphism*. □

Other possibilities would be generated by allowing homomorphisms to be partial functions.

Exercise 2.7.32. Consider a partial Σ -algebra A and its underlying total Σ -algebra A_\perp . Given any Σ -congruence \equiv on A_\perp , removing all pairs involving \perp yields a *strong Σ -congruence on A* . Check that such strong congruences are exactly kernels of strong Σ -homomorphisms; cf. Exercises 1.3.14 and 1.3.18. Check that strong congruences are equivalence relations that preserve and reflect definedness of operations and are closed under defined operations. Kernels of weak Σ -homomorphisms are *weak Σ -congruences*: equivalence relations that are closed under defined operations. Spell out these definitions in detail. For any partial Σ -algebra A and weak Σ -congruence \equiv on A , generalise Definition 1.3.15 to define the *quotient of A by \equiv* , written A/\equiv . Note that an operation is defined in A/\equiv on a tuple of equivalence classes provided that in A it is defined on at least one tuple of their respective elements. Check which of Exercises 1.3.18–1.3.23 carry over. □

Definition 2.7.33 (Term evaluation). Let X be an S -sorted set of variables, let A be a partial Σ -algebra, and let $v: X \rightarrow |A|$ be a (total) S -sorted function assigning values in A to variables in X . Since $|A| \subseteq |A_\perp|$, this is an S -sorted function $v_\perp: X \rightarrow |A_\perp|$, and by Fact 1.4.4 there is a unique (ordinary) Σ -homomorphism $v_\perp^\#: T_\Sigma(X) \rightarrow A_\perp$ which extends v_\perp . Let $s \in S$ and let $t \in |T_\Sigma(X)|_s$ be a Σ -term of sort s ; the *value of t in A under the valuation v* is $v_\perp^\#(t)$ if $v_\perp^\#(t) \neq \perp_s$, and is undefined otherwise. □

Satisfaction of an equation $\forall X \bullet t = t'$, where the values of t and/or t' may be undefined, can be defined in several different ways. Following [BW82b], we use *strong equality* (also known as *Kleene equality*), whereby the equality holds if (for any assignment of values to variables) the values of t and t' either are both defined and equal, or are both undefined. The usual interpretation of definitional equations in recursive function definitions (see for instance Example 4.1.25 and Exercise 4.1.30 below) makes them hold as strong equations. An alternative is *existential equality* (where $=$ is usually written $\stackrel{e}{=}$), whereby the equality holds only when the values of t and t' are defined and equal. When strong equality is used, there is a need for an additional form of axiom called a *definedness formula*: $\forall X \bullet \text{def}(t)$ holds if for any assignment of values to variables, the value of t is defined. These are superfluous with existential equality since $\forall X \bullet \text{def}(t)$ holds iff $\forall X \bullet t \stackrel{e}{=} t$ holds. Definedness formulae with $X = \emptyset$ are called *ground* and are often written without quantification as $\text{def}(t)$.

Exercise 2.7.34. Formalise the definitions of satisfaction of equations (using strong equality) and of definedness formulae. □

Using both equations and definedness formulae as axioms, the definitions of presentation, model of a presentation, semantic consequence, isomorphism, and initial model (with respect to *weak* homomorphisms) are analogous to those given earlier.

Exercise 2.7.35. Spell out the details of these definitions. Note though that not all of the properties of these notions carry over from the standard algebraic framework; for instance, a (weak) bijective homomorphism need not be an isomorphism of partial algebras. \square

Theorem 2.7.36 (Initial model theorem). Any presentation $\langle \Sigma, \mathcal{E} \rangle$ has an initial model I , characterised by the following properties:

- I contains no junk;
- I is minimally defined, i.e. for all $t \in |T_\Sigma|$, t_1 is defined only if $\mathcal{E} \models_\Sigma \text{def}(t)$; and
- I contains no confusion, i.e. for all $t, t' \in |T_\Sigma|_s, s \in S$, t_1 and t'_1 are defined and equal only if $\mathcal{E} \models_\Sigma t = t'$.

Proof sketch. Let Σ_\perp be the signature obtained by adding a constant $\perp_s : s$ to Σ for each sort $s \in S$. Define a congruence $\sim \subseteq |T_{\Sigma_\perp}| \times |T_{\Sigma_\perp}|$ as follows: for $t_1, t_2 \in |T_{\Sigma_\perp}|_s$ for some $s \in S$, $t_1 \sim t_2$ iff any of the following conditions holds:

1. t_1 contains $\perp_{s'}$ and t_2 contains $\perp_{s''}$ for some $s', s'' \in S$;
2. t_1 contains $\perp_{s'}$ for some $s' \in S$, $t_2 \in |T_\Sigma|_s$ (so t_2 does not contain $\perp_{s''}$ for any $s'' \in S$) and $\mathcal{E} \not\models \text{def}(t_2)$, or vice versa;
3. $t_1, t_2 \in |T_\Sigma|_s$, and either $\mathcal{E} \not\models \text{def}(t_1)$ and $\mathcal{E} \not\models \text{def}(t_2)$ or $\mathcal{E} \models t_1 = t_2$.

I is constructed by taking the quotient of T_{Σ_\perp} by \sim , and then regarding congruence classes containing the constants \perp_s as undefined values. \square

Exercise 2.7.37. Complete the above proof by showing that

- \sim is a congruence on T_{Σ_\perp} ;
- $I \models \mathcal{E}$;
- I is an initial model of $\langle \Sigma, \mathcal{E} \rangle$; and
- I has the properties promised in Theorem 2.7.36.

Show that any model of $\langle \Sigma, \mathcal{E} \rangle$ satisfying the properties in Theorem 2.7.36 is isomorphic to I and is therefore an initial model of $\langle \Sigma, \mathcal{E} \rangle$. \square

Exercise 2.7.38. Suppose that we modify Theorem 2.7.36 by replacing the phrase “ t_1 and t'_1 are defined and equal” with “ $I \models_\Sigma t = t'$ ”. Give a counterexample showing that this version of the theorem is false. \square

Exercise 2.7.39. A partial Σ -algebra $A \in \text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$ is a *strongly initial model* of $\langle \Sigma, \mathcal{E} \rangle$ if for every minimally defined $B \in \text{Mod}[\langle \Sigma, \mathcal{E} \rangle]$ containing no junk, there is a unique strong Σ -homomorphism $h: A \rightarrow B$. Show that I is an initial model of $\langle \Sigma, \mathcal{E} \rangle$ iff I is a strongly initial model of $\langle \Sigma, \mathcal{E} \rangle$. \square

Again, reachable and final semantics are applicable for partial algebras as well as initial semantics, and the key points of the standard framework carry over with appropriate changes (for instance, the equational calculus must be modified to deal with definedness formulae as well as equations).

Example 2.7.16 (revisited). Here is a version of the specification NATPRED in which pred is specified to be a partial function:

```

spec NATPRED = sorts Nat
                ops  0: Nat
                    succ: Nat → Nat
                    pred: Nat → Nat
                    ++: Nat × Nat → Nat
                    ××: Nat × Nat → Nat
                ∀m, n: Nat
                • def(0)
                • def(succ(n))
                • pred(succ(n)) = n
                • 0 + n = n
                • succ(m) + n = succ(m + n)
                • 0 × n = 0
                • succ(m) × n = (m × n) + n

```

In initial models of NATPRED , all operations behave as expected, and all are total except for pred , which is undefined only on 0.

Exercise. Show that $\forall m, n: \text{Nat} \bullet \text{def}(m + n)$ and $\forall m, n: \text{Nat} \bullet \text{def}(m \times n)$ are consequences of the definedness axioms for 0 and succ and the equations defining $+$ and \times in reachable models of NATPRED . You will need to use induction, so first formulate an appropriate induction rule scheme and convince yourself that it is sound.

Exercise. Suppose that the axiom $\text{def}(0)$ were removed from NATPRED . Describe the initial models of the resulting presentation. \square

2.7.5 Partial functions: order-sorted algebras

Any partial function amounts to a total function on a restricted domain. The idea of *order-sorted algebra* is to avoid partial functions by enabling the domain of each function to be specified exactly. This is done by introducing *subsorts*, which correspond to subsets at the level of values, and requiring operations to behave in an appropriate fashion when applied to a value of a subsort or when expected to deliver a value of a supersort. A number of different approaches to order-sorted algebra have been proposed, and their relative merits are a matter for debate. Here we follow the approach of [GM92].

Definition 2.7.40 (Order-sorted signature). An *order-sorted signature* is a triple $\Sigma = \langle S, \leq, \Omega \rangle$ where $\langle S, \Omega \rangle$ is an ordinary signature and \leq is a partial order on the set S of sort names, such that whenever $f: s_1 \times \cdots \times s_n \rightarrow s$ and $f: s'_1 \times \cdots \times s'_n \rightarrow s'$ are operations (having the same name and same number of arguments) in Ω and $s_i \leq s'_i$ for all $1 \leq i \leq n$, then $s \leq s'$. When $s \leq s'$ for $s, s' \in S$, we say that s is a

subsort of s' (or equivalently, s' is a *supersort* of s). The subsort ordering is extended to sequences of sorts of equal length in the usual way: $s_1 \dots s_n \leq s'_1 \dots s'_n$ if $s_i \leq s'_i$ for all $1 \leq i \leq n$. \square

The restriction on Ω ([GM92] calls this condition *monotonicity*) is a fairly natural one, keeping in mind that the subsort ordering corresponds to subset on the value level: restricting a function to a subset of its domain may diminish, but not enlarge, its codomain. Note that an effect of this restriction is to rule out overloaded constants.

Throughout the rest of this section, let $\Sigma = \langle S, \leq, \Omega \rangle$ be an order-sorted signature, and let $\widehat{\Sigma} = \langle S, \Omega \rangle$ be the (ordinary) signature underlying Σ .

Definition 2.7.41 (Order-sorted algebra). An *order-sorted Σ -algebra* A is an ordinary $\widehat{\Sigma}$ -algebra, such that:

- for all $s \leq s'$ in Σ , $|A|_s \subseteq |A|_{s'}$; and
- whenever $f: s_1 \times \dots \times s_n \rightarrow s$ and $f: s'_1 \times \dots \times s'_n \rightarrow s'$ are operations (having the same name and same number of arguments) in Ω and $s_1 \dots s_n \leq s'_1 \dots s'_n$, the function $(f: s_1 \times \dots \times s_n \rightarrow s)_A: |A|_{s_1} \times \dots \times |A|_{s_n} \rightarrow |A|_s$ is the set-theoretic restriction of the function $(f: s'_1 \times \dots \times s'_n \rightarrow s')_A: |A|_{s'_1} \times \dots \times |A|_{s'_n} \rightarrow |A|_{s'}$. \square

An effect of the second restriction ([GM92] calls this condition *monotonicity* as well) is to prevent ambiguity in the evaluation of terms; see below.

Definition 2.7.42 (Order-sorted homomorphism). Let A and B be order-sorted Σ -algebras. An *order-sorted Σ -homomorphism* $h: A \rightarrow B$ is an ordinary $\widehat{\Sigma}$ -homomorphism such that $h_s(a) = h_{s'}(a)$ for all $a \in |A|_s$ whenever $s \leq s'$. When h has an inverse, it is an *order-sorted Σ -isomorphism* and we write $A \cong B$. \square

Let X be an S -sorted set (of variables) such that X_s and $X_{s'}$ are disjoint for $s \neq s'$.

Definition 2.7.43 (Order-sorted term algebra). The *order-sorted Σ -algebra $T_\Sigma(X)$ of terms with variables X* is just like $T_{\widehat{\Sigma}}(X)$, except that for any term $t \in |T_\Sigma(X)|_s$ such that $s \leq s'$, we also have $t \in |T_\Sigma(X)|_{s'}$. Let $T_\Sigma = T_\Sigma(\emptyset)$. \square

Exercise 2.7.44. Check that $T_\Sigma(X)$ is an order-sorted Σ -algebra. \square

Example 2.7.45. One way of reformulating NATPRED as an order-sorted specification (see below) will involve introducing a sort $NzNat$ (non-zero natural numbers) such that $NzNat \leq Nat$, with operations $0: Nat$ and $succ: Nat \rightarrow NzNat$. According to the definition of order-sorted term algebra, the term $succ(0)$ has sort Nat as well as $NzNat$, which means that $succ(succ(0))$ is well formed (and has sort Nat as well as $NzNat$). \square

As the above example demonstrates, a given term may appear in more than one carrier of $T_\Sigma(X)$. The following condition on Σ ensures that this does not lead to ambiguity.

Definition 2.7.46 (Regular order-sorted signature). Σ is *regular* if for any $f: s_1 \times \dots \times s_n \rightarrow s$ in Σ and $s'_1 \dots s'_n \leq s_1 \dots s_n$, there is a least $s^*_1 \dots s^*_n$ such that $s'_1 \dots s'_n \leq s^*_1 \dots s^*_n$ and $f: s^*_1 \times \dots \times s^*_n \rightarrow s^*$ is in Σ . \square

Theorem 2.7.47 (Terms have least sorts). *If Σ is regular, then for every term $t \in |T_\Sigma(X)|$ there is a least sort $s \in S$, written $\text{sort}(t)$, such that $t \in |T_\Sigma(X)|_s$.* \square

Exercise 2.7.48. Prove Theorem 2.7.47. What happens when X is an arbitrary S -sorted set, i.e. if we remove the restriction that X_s and $X_{s'}$ are disjoint for $s \neq s'$? \square

Now the definition of term evaluation is analogous to the usual one.

Fact 2.7.49. *Suppose that Σ is regular. Then, for any order-sorted Σ -algebra A and S -sorted function $v: X \rightarrow |A|$, there is exactly one order-sorted Σ -homomorphism $v^\#: T_\Sigma(X) \rightarrow A$ which extends v , i.e. such that $v_s^\#(x) = v_s(x)$ for all $s \in S, x \in X_s$.* \square

Exercise 2.7.50. Define term evaluation. \square

Definition 2.7.51 (Order-sorted equation; satisfaction). Suppose that Σ is regular, and let the equivalence relation \equiv be the symmetric transitive closure of \leq . Order-sorted Σ -equations $\forall X \bullet t = t'$ are as usual, except that we require $\text{sort}(t) \equiv \text{sort}(t')$ (in other words, $\text{sort}(t)$ and $\text{sort}(t')$ are in the same connected component of $\langle S, \leq \rangle$) instead of $\text{sort}(t) = \text{sort}(t')$. An order-sorted Σ -algebra A satisfies an order-sorted Σ -equation $\forall X \bullet t = t'$, written $A \models_\Sigma \forall X \bullet t = t'$, if the value of t in $|A|_{\text{sort}(t)}$ and the value of t' in $|A|_{\text{sort}(t')}$ coincide for every S -sorted function $v: X \rightarrow |A|$. \square

A problem with this definition is that satisfaction of order-sorted Σ -equations is not preserved by order-sorted Σ -isomorphisms (compare Exercise 2.1.5). The following condition on Σ ensures that this anomaly does not arise.

Definition 2.7.52 (Coherent order-sorted signature). $\langle S, \leq \rangle$ is filtered if for any $s, s' \in S$ there is some $s'' \in S$ such that $s \leq s''$ and $s' \leq s''$. $\langle S, \leq \rangle$ is locally filtered if each of its connected components is filtered. Σ is coherent if $\langle S, \leq \rangle$ is locally filtered and Σ is regular. \square

Exercise 2.7.53. Find Σ, A, B and e such that Σ is regular, $A \models_\Sigma e$ and $A \cong B$ but $B \not\models_\Sigma e$. Show that if Σ is coherent then this is impossible. \square

The definitions of order-sorted presentation, model of an order-sorted presentation, semantic consequence, and initial model are analogous to those given earlier. For every order-sorted presentation $\langle \Sigma, \mathcal{E} \rangle$ such that Σ is coherent, an initial model may be constructed as a quotient of T_Σ [GM92]. There is a version of the equational calculus that is sound and complete for coherent signatures [GM92], and the use of term rewriting for proof as discussed in Section 2.6 is sound, provided that each rewrite rule $t \rightarrow t'$ is sort decreasing, i.e. $\text{sort}(t') \leq \text{sort}(t)$ [KKM88].

Example 2.7.16 (revisited). Here is a version of the specification NATPRED in which *pred* is specified to be a total function on the non-zero natural numbers:

```

spec NATPRED = sorts NzNat ≤ Nat
ops 0: Nat
      succ: Nat → NzNat
      pred: NzNat → Nat
      ++: Nat × Nat → Nat
      ××: Nat × Nat → Nat
      ∀m, n: Nat
        • pred(succ(n)) = n
        • 0 + n = n
        • succ(m) + n = succ(m + n)
        • 0 × n = 0
        • succ(m) × n = (m × n) + n

```

In this version of NATPRED, there are terms that are not well formed in spite of the fact that each operator application seems to be to a value in its domain. For example, consider the following “term”:

$$\text{pred}(\text{succ}(0) + \text{succ}(0)).$$

According to the signature of NATPRED, $\text{succ}(0) + \text{succ}(0)$ is a term of sort *Nat*; it is not a term of sort *NzNat* in spite of the fact that its value is non-zero. In the term algebra, *pred* applies only to terms of sort *NzNat*; thus the application of *pred* to $\text{succ}(0) + \text{succ}(0)$ is not defined. One way of getting around this problem might be to add additional operators to the signature of NATPRED:

```

spec NATPRED = sorts NzNat ≤ Nat
ops ...
      ++: NzNat × Nat → NzNat
      ++: Nat × NzNat → NzNat
      ××: NzNat × NzNat → NzNat
      ...

```

Then $\text{succ}(0) + \text{succ}(0)$ is a term of sort *NzNat*, as desired. Unfortunately, this signature is not regular. (**Exercise:** Why not? What can be done to make it regular?)

An alternative is to use a so-called *retract*, an additional operation for converting from a sort to one of its subsorts:

```

spec NATPRED = sorts NzNat ≤ Nat
ops ...
      r: Nat → NzNat
      ∀m, n: Nat, k: NzNat
        • ...
        • r(k) = k

```

Now, the term $\text{pred}(r(\text{succ}(0) + \text{succ}(0)))$ is well formed, and is equal to $\text{succ}(0)$ in all models of NATPRED. In the words of [GM92], inserting the retract *r* into $\text{pred}(r(\text{succ}(0) + \text{succ}(0)))$ gives it “the benefit of the doubt”, and the term is “vindicated” by the fact that it is equal to a term that does not contain *r*. The term

$pred(r(0))$ is also well formed, but in the initial model of $NATPRED$ this term is equal only to other terms containing the retract r , and can thus be regarded as an error message. The use of retracts (which can be inserted automatically) is well behaved under certain conditions on order-sorted presentations [GM92].

Another version of $NATPRED$ is obtained by using an *error supersort* for the codomain of $pred$ rather than a subsort for its domain:

spec $NATPRED = \mathbf{sorts}$ $Nat \leq Nat?$
ops $0: Nat$
 $succ: Nat \rightarrow Nat$
 $pred: Nat \rightarrow Nat?$
 $-- + --: Nat \times Nat \rightarrow Nat$
 $-- \times --: Nat \times Nat \rightarrow Nat$
 $\forall m, n: Nat$

- $pred(succ(n)) = n$
- $0 + n = n$
- $succ(m) + n = succ(m + n)$
- $0 \times n = 0$
- $succ(m) \times n = (m \times n) + n$

The sort $Nat?$ may be thought of as Nat extended by the addition of an error value corresponding to $pred(0)$.

Here we have the same problem with ill-formed terms as before; an example is the term $succ(pred(succ(0)))$. Again, retracts solve the problem. In this case, the required retract is the operation $r: Nat? \rightarrow Nat$, defined by the axiom $\forall n: Nat \bullet r(n) = n$. □

Exercise 2.7.54. Try to view the error algebra approach presented in Section 2.7.3 as a special case of order-sorted algebra. □

2.7.6 Other options

The previous sections have mentioned only a few of the ways in which the standard framework can be improved to make it more suitable for particular kinds of applications. A great many other variations are possible; a few of these are sketched below.

Example 2.7.55 (First-order predicate logic). Signatures may be modified to enable them to include (typed) *predicate names* in addition to operation names, e.g. $-- \leq --: Nat \times Nat$. Atomic formulae are then formed by applying predicates to terms; in *first-order predicate logic with equality*, the predicate $-- = --: s \times s$ is implicitly available for any sort s . Formulae are built from atomic formulae using logical connectives and quantifiers. Algebras are modified to include relations on their carriers to interpret predicate names; the interpretation of the built-in equality predicate (if available) may be forced to be the underlying equality on values,

or it may merely be required to be a congruence relation. Homomorphisms are required to respect predicates as well as operations. The satisfaction of a *sentence* (a formula without free variables) by an algebra is as in first-order logic. See Example 4.1.12 for details of the version of first-order predicate logic with equality we will use. Presentations involving predicates and first-order axioms are appropriate for the specification of programs in *logic programming languages* such as Prolog, where the Horn clause fragment of first-order logic is used for writing the programs themselves. Note that such presentations may have no models at all, but even if they have some models, they may have no initial models (see Example 2.7.11) or no final models (see Exercise 2.7.14), or even no reachable models. (**Exercise:** Give a specification with first-order axioms having some models but no reachable model.) \square

Example 2.7.56 (Higher-order functions). Higher-order functions (taking functions as parameters and/or returning functions as results) can be accommodated by interpreting certain sort names as (subsets of) function spaces. Given a set S of (base) sorts, let S^{\rightarrow} be the closure of S under formation of function types: S^{\rightarrow} is the smallest set such that $S \subseteq S^{\rightarrow}$ and for all $s_1, \dots, s_n, s \in S^{\rightarrow}$, $s_1 \times \dots \times s_n \rightarrow s \in S^{\rightarrow}$. Then a *higher-order signature* Σ is a pair $\langle S, \Omega \rangle$ where Ω is an S^{\rightarrow} -indexed set of operation names. This determines an ordinary signature Σ^{\rightarrow} comprised of the sort names S^{\rightarrow} and the operation names in Ω together with operation names $apply: (s_1 \times \dots \times s_n \rightarrow s) \times s_1 \times \dots \times s_n \rightarrow s$ for every $s_1, \dots, s_n, s \in S^{\rightarrow}$. Note that, except for the various instances of *apply*, all the operations in Σ^{\rightarrow} are constants, albeit possibly of “functional” sort. A *higher-order Σ -algebra* is just an ordinary (total) Σ^{\rightarrow} -algebra, and analogously for the definitions of higher-order Σ -homomorphism, reachable higher-order Σ -algebra, higher-order Σ -term, higher-order Σ -equation, satisfaction of a higher-order Σ -equation by a higher-order Σ -algebra, and higher-order presentation. A higher-order Σ -algebra A is *extensional* if for all sorts $s_1 \times \dots \times s_n \rightarrow s \in S^{\rightarrow}$ and values $f, g \in |A|_{s_1 \times \dots \times s_n \rightarrow s}$, $f = g$ whenever $apply_A(f, a_1, \dots, a_n) = apply_A(g, a_1, \dots, a_n)$ for all $a_1 \in |A|_{s_1}, \dots, a_n \in |A|_{s_n}$. Any extensional higher-order algebra is isomorphic to an (extensional) algebra A , where every carrier $|A|_{s_1 \times \dots \times s_n \rightarrow s}$ is a subset of the function space $|A|_{s_1} \times \dots \times |A|_{s_n} \rightarrow |A|_s$ and all the operations $apply_A$ are the usual function application. A higher-order Σ -algebra A is a *model* of a presentation $\langle \Sigma, \mathcal{E} \rangle$ if $A \models_{\Sigma} \mathcal{E}$, A is extensional, and A is reachable. The reachability requirement (no junk) means that $|A|_{s_1 \times \dots \times s_n \rightarrow s}$ will almost never be the full function space $|A|_{s_1} \times \dots \times |A|_{s_n} \rightarrow |A|_s$: only the functions that are denotable by ground terms will be present in $|A|_{s_1 \times \dots \times s_n \rightarrow s}$. Higher-order (equational) presentations always have initial models [MTW88]. \square

Example 2.7.57 (Polymorphic types). Standard ML [Pau96] and some other programming languages define *polymorphic types* such as α list (instances of which include `bool list` and `(bool list) list`) and *polymorphic values* of those types, such as `head: $\forall \alpha. \alpha$ list $\rightarrow \alpha$` (which is then applicable to values of types such as `bool list` and `(bool list) list`, yielding results of types `bool` and `bool list`, respectively). To specify such types and functions, signatures are modified to contain *type constructors* in place of sort names; for example, `list` is a unary type constructor and `bool` is a nullary type constructor. Terms built using these type

constructors and *type variables* (such as α above) are the *polymorphic types* of the signature. The set Ω of operation names is then indexed by non-empty sequences of polymorphic types, where $f \in \Omega_{t_1 \dots t_n, t}$ means $f: \forall FV(t_1) \cup \dots \cup FV(t_n) \cup FV(t) \bullet t_1 \times \dots \times t_n \rightarrow t$. There are various choices for algebras over such signatures. Perhaps the most straightforward choice is to require each algebra A to incorporate a (single-sorted) *algebra of carriers* $\text{Carr}(A)$, having sets interpreting types as values and with an operation to interpret each type constructor. Then, for each operation $f \in \Omega_{t_1 \dots t_n, t}$ and for each instantiation of type variables $i: V \rightarrow |\text{Carr}(A)|$, A has to provide a function $f_{A,i}: i^\#(t_1) \times \dots \times i^\#(t_n) \rightarrow i^\#(t)$. Various conditions may be imposed to ensure that the interpretation of polymorphic operations is *parametric* in the sense of [Str67], by requiring $f_{A,i}$ and $f_{A,i'}$ to be appropriately related for different type variable instantiations i, i' ; see Exercise 3.4.40 for a hint in this direction. Another choice would be to interpret each type as the set of equivalence classes of a *partial equivalence relation* on a model of the untyped λ -calculus [BC88]. Axioms contain (universal) quantifiers for type variables in addition to quantifiers for ordinary variables, as in System F [Gir89]; alternatively, type variable quantification may be left implicit, as in Extended ML [KST97]. \square

Example 2.7.58 (Non-deterministic functions). Non-deterministic functions may be handled by interpreting operation names in algebras as relations, or equivalently as set-valued functions. Homomorphisms are required to preserve possible values of functions: for any homomorphism $h: A \rightarrow B$ and operation $f: s_1 \times \dots \times s_n \rightarrow s$, if a is a possible value of $f_A(a_1, \dots, a_n)$ then $h_s(a)$ is a possible value of $f_B(h_{s_1}(a_1), \dots, h_{s_n}(a_n))$. Universally quantified inclusions between sets of possible values may be used as axioms: $t \subseteq t'$ means that every possible value of t is a possible value of t' . \square

Example 2.7.59 (Recursive definitions). Following [Sco76], partial functions may be specified as least solutions of recursive equations, where “least” is with respect to an ordering on the space of functions of a given type. To accommodate this, we can use *continuous algebras*, i.e. ordinary (total) Σ -algebras with carriers that are complete partially ordered sets (so-called *cpos*) and with operation names interpreted as *continuous functions* on these sets. See Example 3.3.14. The “bottom” element \perp of the carrier for a sort, if it exists, represents the completely undefined value of that sort. The order on carriers induces an order on (continuous) functions in the usual fashion. A homomorphism between continuous algebras is required to be continuous as a function between cpos. It is possible to define a language of axioms that allows direct reference to least upper bounds of chains (see Example 4.1.22), and/or to the order relation itself. Such techniques may also be used to specify infinite data types such as *streams*. \square

2.8 Bibliographical remarks

Much of the material presented here is well known, at least in its single-sorted version, in universal algebra as a branch of mathematics. Standard references are [Grä79] and [Coh65]. We approach this material from the direction of computer science — see [Wec92] and [MT92] — and present the fundamentals of equational specifications as developed in the 1970s ([Zil74], [Gut75], [GTW76]); see also [EM85] for an extended monograph-style presentation.

The simplest and most limited form of a specification is a “bare” signature, and this is what is used to characterise classes of algebras (program modules) in modularisation systems for programming languages — see, e.g., Standard ML [MTHM97], [Pau96], where such characterisations are in fact called signatures, *type classes* in Haskell [Pey03] and *concepts* in C++ [C++09]. Presentations correspond to Extended ML signatures [ST85] and to C++ concepts containing axioms.

The first appearance of the Satisfaction Lemma (Lemma 2.1.8) in the algebraic specification literature was in [BG80], echoing the semantic consequences of the definition of (theory) interpretations in logic [End72]. This fundamental link between syntax and semantics will become one of the cornerstones of later development starting in Chapter 4.

One topic that is only touched upon here (see, e.g., Theorem 2.2.10) is the expressive power of specifications. See [BT87] for a comprehensive survey of what is known about the expressive power of the framework presented in this chapter. The main theorem is the one mentioned at the beginning of Section 2.7.

We make a distinction between presentations and theories that is not present in some other work. This distinction surfaces in the definition of theory morphisms (Definition 2.3.11). For two presentations (not necessarily theories) $\langle \Sigma, \mathcal{E} \rangle$ and $\langle \Sigma', \mathcal{E}' \rangle$, [Gan83] takes a signature morphism $\sigma: \Sigma \rightarrow \Sigma'$ to be a specification morphism $\sigma: \langle \Sigma, \mathcal{E} \rangle \rightarrow \langle \Sigma', \mathcal{E}' \rangle$ if $\sigma(\mathcal{E}) \subseteq \mathcal{E}'$. Such a σ is referred to as an “axiom-preserving theory morphism” in [Mes89]. Exercise 2.3.15 shows that this is not equivalent to our definition of theory morphism between the theories presented by those presentations. Another possibility is to require σ to map only the *ground* equations in \mathcal{E} to equations in $Cl_{\Sigma'}(\mathcal{E}')$, as in [Ehr82]. These alternative definitions seem unsatisfactory since they make little or no sense on the level of models, in contrast to the relationship between theory and model levels for theory morphisms given by Proposition 2.3.13. We will later (Definition 5.5.1) define *specification morphisms*, as a generalisation of morphisms between presentations, relying on this relationship.

The many-sorted equational calculus is presented in [GM85] together with a proof that it is sound and complete. This builds on the standard equational calculus [Bir35], but the modifications needed to deal with empty carriers in the many-sorted context came as a surprise at the time. Our choice of rules in Section 2.4 is different from this standard version but the two systems are equivalent (Exercise 2.4.14) and the proofs of soundness and completeness are analogous.

The initial algebra approach to specification (Section 2.5) is the classical one. It originated with the seminal paper [GTW76], and was further developed by Hartmut Ehrig and his group; see [EM85] for a comprehensive account.

Example 2.5.24 and Exercise 2.5.25 point at useful ways to make inductive proofs easier by providing derived induction rule schemes, as possible, for instance, in the logics of Larch [GH93] and CASL [Mos04] and their proof support systems (LP [GG89] and HETS [MML07], respectively); see also Chapter 6 of [Far92].

The proof of the incompleteness theorem for initial semantics (Theorem 2.5.26) from [MS85] follows [Nou81] where it was used to show that the equational calculus with a specific induction rule scheme is not complete. An alternative to adding induction rules to the equational calculus is to restrict attention to so-called ω -complete presentations; these are presentations $\langle \Sigma, \mathcal{E} \rangle$ for which the equational calculus itself yields all of the Σ -equations that hold in initial models of $\langle \Sigma, \mathcal{E} \rangle$ [Hee86]. Then the problem becomes one of finding an ω -complete presentation corresponding to a given presentation. By the incompleteness theorem, this is not always possible.

There is a substantial body of theory on term rewriting systems; Section 2.6 is only the tip of the iceberg. For much more on the topic, and for the details of the Knuth-Bendix completion algorithm [KB70] that have been omitted in Section 2.6, see [DJ90], [Klo92], [BN98], [Kir99] and [Ter03]. See [KM87] or [DJ90] for a discussion of proof by consistency, which originated with [Mus80]. Like most work in this area, all these restrict attention to the single-sorted case. See [EM85] for a treatment of the many-sorted case, up to the soundness and completeness theorems for conversion, without our simplifying assumption (cf. Exercise 2.6.11).

In the case of reachable and final semantics, it is usual to look at reachable or final *extensions* of algebras (alternative terminology: hierarchical specifications), rather than at the reachable or final interpretation of a completed specification. See [BDP⁺79] or [WB82] for reachable semantics, and [GGM76] or [Wan79] for final semantics. Under appropriate conditions, the reachable models of a presentation form a complete lattice, with the initial model at one extreme and the final model at the other; see [GGM76] and [BWP84]. For such hierarchical specifications, an incompleteness theorem that is even stronger than Theorem 2.5.26 may be proved: no sound proof system can derive all *ground* equational consequences of such specifications; see [MS85].

The first attempt to specify errors by distinguishing error values from OK values was [Gog78]. More details of the approach outlined in Section 2.7.3 can be found in [GDLE84]. The final semantics of error presentations is discussed in [Gog85]. See [BBC86] for an alternative approach which is able to deal with examples like the one discussed in Exercise 2.7.29.

More details of the approach to partial algebras outlined in Section 2.7.4 can be found in [BW82b]. Weak Σ -homomorphisms are called total Σ -homomorphisms there. Alternative approaches to the specification of partial algebras are presented in [Rei87] and [Kre87], and more recently in [Mos04]. See [Bur86] for a comprehensive analysis of the various alternative definitions of the basic notions.

See [GM92], further refined in [Mes09], for more on the approach to order-sorted algebra in Section 2.7.5. Alternative approaches include [Gog84], [Poi90] and [Smo86], which is sometimes referred to as “universal” order-sorted algebra to distinguish it from “overloaded” order-sorted algebra as presented here. A uni-

versal order-sorted algebra contains a single universe of values, where a sort corresponds to a subset of the universe and each operation name identifies a (single) function on the universe. A compromise is in rewriting logic [Mes92] as implemented in Maude [CDE⁺02]. See [Mos93] and [GD94a] for surveys comparing the different approaches. [GD94a] discusses how some of the definitions and results in Section 2.7.5 can be generalised by dropping or weakening the monotonicity requirements on order-sorted signatures and order-sorted algebras. Yet a different approach to subsorting is taken in CASL [Mos04] where subsort coercions may be arbitrary injective functions rather than merely inclusions.

First-order predicate logic has been used as a framework for algebraic specification in various approaches; see for instance CIP-L [BBB⁺85] and CASL [Mos04]. See [Poi86], [MTW88], [Mei92] and [Qia93] for different approaches to the algebraic specification of higher-order functions. Frameworks that cater for the specification of polymorphic types and functions are described in [Mos89], [MSS90] and [KST97]. See [Nip86] for more on algebras with non-deterministic operations; for a different approach using relation algebra, see [BS93]. See [WM97] for a comprehensive overview. Soundness and completeness of term rewriting for non-deterministic specifications is studied in [Hus92]. Continuous algebras and the use of Scott-style domain-theoretic techniques in algebraic specification were first discussed in [GTWW77]. See [Sch86] or [GS90] for much more on domain theory itself. Although these and other extensions to the standard framework have been explored separately, the few attempts that have been made to combine such extensions (see, e.g., [AC89] and [Mos04]) have tended to reveal new problems.



<http://www.springer.com/978-3-642-17335-6>

Foundations of Algebraic Specification and Formal
Software Development

Sannella, D.; Tarlecki, A.

2012, XVI, 584 p., Hardcover

ISBN: 978-3-642-17335-6