

Essential Concepts of Algebraic Specification and Program Development

Donald Sannella¹ and Andrzej Tarlecki²

¹LFCS, Department of Computer Science, University of Edinburgh, Edinburgh, UK

²Institute of Informatics, Warsaw University and Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland

Keywords: Algebraic specification; Formal program development; Specification and program structure; Implementation of specifications; Behavioural equivalence

Abstract. The main ideas underlying work on the model-theoretic foundations of algebraic specification and formal program development are presented in an informal way. An attempt is made to offer an overall view, rather than new results, and to focus on the basic motivation behind the technicalities presented elsewhere.

Introduction

The long-term goal of work on algebraic specification is to provide a formal basis to support the systematic development of correct programs from specifications by means of verified refinement steps. There has been a large body of technical work directed towards this important goal. Many interesting concepts have been introduced and quite a number of non-trivial results have been stated and proved (see [BKL91] for a review and a comprehensive list of references). Instead of providing yet another piece in the puzzle, in this paper we sketch on a rather informal level our views on how some of the existing pieces fit into an overall picture of what is important in the light of the ultimate goal. We focus on the motivations for certain technicalities that we think are of crucial importance, only suggesting, rather than presenting in full detail, the technicalities themselves. A past paper with similar aims is [GHW82].

The literature already mentions many of the points we make here, such as the use of “institutions” to ensure sufficient generality of the proposed framework,

and the use of “constructor implementations” to capture the essence of program development steps (including steps that involve a decomposition into independent programming tasks). Some of these ideas have been hidden amongst the technical definitions and results, and we think they are worth restating here more prominently, with more careful arguments in some cases. For example, we provide a more detailed justification for the use of model classes, rather than theories, as the appropriate semantic domain for specifications. We also give a simple explanation of the somewhat subtle interplay between behavioural equivalence, “stability” and refinement in formal development.

Examples are provided to illustrate some of the points we make. Chosen for simplicity rather than to impress, they are kept as small as possible, and some are contrived just to illustrate a particular point. This should not be taken as an indication that large examples cannot be handled, of course. For the sake of concreteness, and to show how the ideas fit into the context of existing programming languages, in examples we use notation and concepts borrowed from the Standard ML (SML) programming language [Pau91] and Extended ML (EML) specification framework [San91, KST97].

Since the emphasis here is on motivation and intuition, citations to the literature refer the interested reader to papers where complete technical details may be found. These topics are also covered in detail in a forthcoming monograph [SaT9?]. For readers who prefer even less detail, the material in Sections 3.1, 8.1 and 9.1 may be skipped on first reading.

1. The Logical Framework

The overall aim of work on algebraic specification is to provide semantic foundations for the development of programs that are *correct* with respect to their requirements specifications. In other words, the program developed must exhibit the required input/output behaviour. We view the correctness of a program as its most crucial property. Other desirable properties (efficiency, robustness, reliability etc.) are disregarded in this work. Of course, this does not mean that these properties are unimportant, but this approach does not provide any formal means for their analysis.

The assumption that the correctness of the input/output behaviour of a program takes precedence over all its other properties allows us to abstract away from concrete details of code and algorithms, and to model program functions as mathematical functions. Such functions are never considered in isolation, but always in units (program modules) comprising a collection of related functions together with the data domains they operate on. At this level of abstraction we are dealing directly with the information essential for the analysis of program correctness, without the burden of irrelevant details. This leads to the most fundamental assumption underlying work on algebraic specification: programs are modelled as many-sorted *algebras*. This assumption fits most directly into the *functional programming* paradigm, but there is a natural way of generalizing it to handle e.g. imperative programs; see below.

We refrain from recalling the formal definition of many-sorted algebra (see e.g. [EhM85]). It is enough to know that an algebra consists of a collection of *carriers* (sets of data) and *operations* on them. Algebras are classified by *signatures*, naming the algebra components (sorts and operations) and thus providing the basic vocabulary for using the program and for making assertions about its

properties. The class of all Σ -algebras (algebras over the signature Σ) will be denoted by $Alg(\Sigma)$. For any program P , the algebra it denotes is written as $\llbracket P \rrbracket \in Alg(Sig(P))$, where $Sig(P)$ is the underlying signature of P .

For any signature, we need a logical system for describing properties of algebras over that signature. Many-sorted equational logic (cf. [GoM85, EhM85]) is the most commonly-used system for this purpose, at least in the area of algebraic specification. Properties of Σ -algebras (or rather, of their operations) may be described by universally-quantified *equations* over Σ , via the definition of what it means for a Σ -algebra A to *satisfy* a Σ -equation φ , written $A \models \varphi$. This also determines a notion of *logical consequence*: a set of equations Φ entails an equation φ , written $\Phi \models \varphi$, if every algebra that satisfies all the equations in Φ also satisfies φ . Here is a simple example, where a signature is accompanied by a list of equational *axioms*, presented using a hopefully self-explanatory notation:

sorts	$nat, list$
opns	$0 : nat$ $succ : nat \rightarrow nat$ $nil : list$ $cons : nat \times list \rightarrow list$ $head : list \rightarrow nat$ $tail : list \rightarrow list$
axioms	$\forall x:nat. \forall l:list. head(cons(x, l)) = x$ $\forall x:nat. \forall l:list. tail(cons(x, l)) = l$

For example now:

$$\left\{ \begin{array}{l} \forall x:nat. \forall l:list. head(cons(x, l)) = x, \\ \forall x:nat. \forall l:list. tail(cons(x, l)) = l \end{array} \right\} \models \forall x, y:nat. head(tail(cons(x, cons(y, nil)))) = y$$

Very rarely in the process of program development does the user work with just a single signature: operations and sorts of data are renamed, added and hidden as the need arises. To take account of this, signatures are equipped with a notion of *signature morphism* (cf. [EhM85]). A signature morphism $\sigma : \Sigma \rightarrow \Sigma'$ maps the sorts and operations of Σ to those of Σ' . This determines in a natural way a translation of any Σ -equation φ to a Σ' -equation $\sigma(\varphi)$, and on the semantic level, a translation of any Σ' -algebra $A' \in Alg(\Sigma')$ to its *reduct* $A'|_{\sigma} \in Alg(\Sigma)$ — notice the change of direction! A typical case is when $\sigma : \Sigma \rightarrow \Sigma'$ is a signature inclusion; then $A'|_{\sigma}$, written $A'|_{\Sigma}$ in this case, is just A' with the interpretation of symbols not in Σ removed. A crucial property is that these two translations are compatible with satisfaction: for any Σ -equation φ and Σ' -algebra A' , $A'|_{\sigma} \models \varphi$ iff $A' \models \sigma(\varphi)$.

The above framework is often criticised (quite rightly!) as rather restrictive and cumbersome to use in practice. Some important features of programs, for example non-termination and higher-order functions, are difficult to model in algebras; equations are not expressive enough to conveniently capture certain properties that one may want to state as requirements. For instance, one may wish to add to the list of axioms above the following two properties which cannot be expressed in equational logic and which are stated here as sentences in first-order logic with equality (using the standard notation for negations of equalities):

axioms $\forall n:\text{nat}. \text{succ}(n) \neq 0$
 $\forall x:\text{nat}. \forall l:\text{list}. \text{cons}(x, l) \neq \text{nil}$

Fortunately, this deficiency is relatively easy to overcome using the concept of *institution*. This concept was introduced by Goguen and Burstall [GoB84] to capture the informal notion of logical system and was strongly influenced by the understanding of this notion in the theory of specifications (see [Bar74] for an early account of *abstract model theory* covering similar ideas approached from the viewpoint of classical logic and model theory, and [BaF85] for a compendium of more recent work in this area).

An institution defines a notion of signature together with for any signature Σ , a set of Σ -sentences, a class of Σ -models and a satisfaction relation between Σ -models and Σ -sentences. Moreover, signatures come equipped with a notion of signature morphism. Any signature morphism induces a translation of sentences and a translation of models (the latter going in the opposite direction as above). The only semantic requirement is that when we change signatures using a signature morphism, the induced translations of sentences and of models preserve the satisfaction relation. Many standard logical systems have been presented explicitly as institutions, see e.g. [GoB92] and [SaT9?]. These include first-order predicate logic with and without equality, and logical systems for specifying partial functions, exception handling, and simple imperative programs. (Most of the examples in this paper are couched in first-order predicate logic with equality.) It should be easy to see that any usual logical system with a well-defined model theory fits into this mould.

Everything below, barring concrete examples, works in the framework of an arbitrary institution, even though for the reader's convenience we avoid "institutional jargon" and refer to "algebras" rather than "models" in the sequel. Consequently, everything in this paper applies to many different concepts of "signature", "algebra" and "sentence" used in the theory and practice of software specification. This point of view gives rise to "reusable" methodologies, theorems, and (ultimately) tools, all of which can be seen as parameterized by an arbitrary institution. See e.g. [BeV87, SaT87, SaT88a, SaT88b, Far92, SST92, DGS93] for work on various aspects of software specification and development that is generic in this sense. Other formulations of general logic have been used for similar purposes, see e.g. [FiS88] and [EBO93].

Strict followers of the early approaches to algebraic specification might view this generalization as an alarming departure, and might protest that this is not algebraic specification at all. In our view the essential idea of algebraic specification is the stress on "algebra-like" models and the use of logical axioms to describe such models. The use of ordinary many-sorted algebras and equations is but a special case of this. Just as it was necessary to generalize from classical single-sorted algebras to many-sorted algebras in order to deal with programs handling several kinds of data, it is necessary to adopt more complicated models to deal with other features of programming languages (polymorphism, higher-order functions, infinite behaviour, updateable references, lazy evaluation, etc.). The essence is that we need a notion of semantic structure that is detailed enough to capture the program properties we want to analyse and abstract enough to make this analysis feasible. Moreover, to specify and reason about programs, we need a logical system with a model theory based on such structures; again, there is a tradeoff here between expressive power and ease of use. There seems to be no single kind of semantic structure that suffices for all purposes, and different

logical systems are appropriate for the analysis of different facets of program behaviour. The multiplicity of logical frameworks seems to be a natural state of affairs rather than indicative of a failure to find the right approach.

2. Specifications

What is a specification? Clearly, since our aim is a formal approach to software development, specifications must be objects as formal as (for example) programs are. That is, we have to have a formal language to write specifications down and to provide a vehicle on which formal techniques to manipulate specifications may be based. It is important for such a *specification language* to provide a collection of convenient notational conventions that are easy to understand and use. One of the basic constituents of a specification will be a list of axioms the specified program is required to satisfy.

A specification formalism must offer means for building complex *structured* specifications by combining and extending simpler ones. A specification of a real-life system typically states a huge number of properties, and building such a specification in an unstructured, monolithic way would result in a long list of axioms which would be neither understandable nor useful. Moreover, the structure of a specification may be used to express intangible aspects of the specifier's knowledge of the problem, such as the degree to which the entities and concepts described in the specification are interrelated. For this purpose, a specification language must provide some *specification-building operations* used to put together small specifications to form more complex ones [BuG77, BuG80]. Then, an understanding of a large specification is achieved via an understanding of its components. This is the principle of *compositionality*: the meaning of a composite object depends only on the meanings of its immediate sub-components.

Various other activities involving specifications can exploit their structure. For example, proofs of consequences of a specification can be usefully guided by its structure [SaB83, HST94] (cf. Section 5). But this principle must not be taken too far: for example, the structure of a specification should not constrain the final structure of its implementations. This is one of the consequences of the famous dogma that a specification should describe only the *whats* of the specified software without constraining any of its *hows*. Requiring the structure of the initial specification to be preserved in its implementation would be unrealistic and unreasonable, even though this has been explicitly suggested by some (e.g. [GoB80, MoA91]) and is implicit in the approaches taken by others. The aims of structuring requirements specifications are often contradictory with the aims of structuring software. See for instance [FiJ90] for a nice discussion of a practical example where such a discrepancy occurs. Section 6 gives a simple example illustrating this point, and Section 7 indicates how the design of the structure of an implementation may be brought into our framework.

Choosing appropriate specification-building operations to be included in a specification language is a non-trivial task, even though most specification languages share certain common operations such as those given below. The choice involves a certain trade-off between the expressive power of the specification language and the ease of understanding and dealing with the operations. One way to circumvent this problem is to first develop a *kernel* language consisting of a minimal set of very powerful, but perhaps awkward to use operations, and then build on top of it a higher-level, more user-friendly language, perhaps sacrificing

some of the expressive power to achieve ease of use and ease of understanding. Such an approach has been taken with the ASL kernel specification language [SaW83, Wir86, SaT88a], on top of which languages such as PLUSS [BGM89] and Extended ML [SaT86] have been built.

In this paper we will neither present nor use a full-blown specification language. In examples we will rely only on the following three simple specification-building operations:

Basic specifications: The specification

sorts S
opns Ω
axioms Φ

describes algebras over the signature with sorts S and operations Ω that satisfy the axioms Φ .

Enrichment: The specification

enrich SP **by** **sorts** S
 opns Ω
 axioms Φ

describes algebras that add the sorts S and operations Ω to algebras described by SP in such a way that the axioms Φ are satisfied.

Hiding: The specification

hide **sorts** S
 opns Ω
in SP

describes those algebras obtained by removing the sorts S and operations Ω from algebras described by SP .

In examples we will omit keywords like **sorts** when the corresponding list of items is empty.

Example 1. Here are some simple specifications:

$BOOL =$ **sorts** $bool$
 opns $true : bool$
 $false : bool$
 axioms $true \neq false$
 $\forall x:bool. x = true \vee x = false$

$INT =$ **enrich** $BOOL$ **by**
 sorts int
 opns $0 : int$
 $succ : int \rightarrow int$
 $pred : int \rightarrow int$
 axioms \dots induction scheme for $int \dots$
 $\forall x:int. pred(x) \neq x \wedge succ(x) \neq x$
 $\forall x:int. pred(succ(x)) = x \wedge succ(pred(x)) = x$

We're glossing over the details of induction schemes here — think of each of these as either an infinite set of first-order axioms given by the usual elementary induction scheme, or (not equivalently!) as a single second-order axiom, or a single infinitary disjunction. An alternative which is more usual in specification languages is to introduce a separate specification-building operation that restricts the class of admissible realizations of a specification to reachable algebras only, see e.g. [SaW83, Wir86] (this is equivalent to the additional second-order axiom or infinitary disjunction). Yet another potential possibility is to restrict the class of algebras considered to reachable algebras from the very beginning [BaW82]. Also note that the axioms of *INTLIST* do not constrain the value of *head(nil)* or *tail(nil)*, meaning that any result is acceptable. An alternative is to specify some error behaviour — see [BKL91] for various approaches.

$SORT =$ **hide opns** *is_sorted* in $SORT1$

The use of the hiding operation in *SORT* means that the *is_sorted* operation does not appear in algebras described by *SORT*. It appears in the specification as an auxiliary operation which allows us to formulate the axioms for *sort* conveniently. (Don't confuse the specification *SORT* containing the operation *sort* with the noun "sort" and the keyword **sorts**!) The observant reader might have noticed that the axioms of *SORT1* do not require *sort* to preserve repetitions in its input. We exploit this to illustrate some further points in Sections 4, 6 and 7.

The above does not take into account the fact that typical programming languages like SML provide booleans, integers and lists as built-in types. Rather than re-specifying and then re-implementing them from scratch, we could follow Extended ML [KST94, KST97] by assuming that all programs implicitly extend all the built-in types and values available, so programs and specifications may freely refer to them. Modifying the above specifications along these lines, assuming that the built-in types and values of SML are available, yields the following:

```

INTORD =
  opns    po : int × int → bool
  axioms  ∀x:int. po(x, x) = true
          ∀x, y:int. po(x, y) = true ∧ po(y, x) = true ⇒ x = y
          ∀x, y, z:int. po(x, y) = true ∧ po(y, z) = true ⇒
                                po(x, z) = true

INTLIST = enrich INTORD by
  opns    head : int list → int
          tail : int list → int list
          is_in : int × int list → bool
  axioms  ∀x:int. ∀l:int list. head(x::l) = x
          ∀x:int. ∀l:int list. tail(x::l) = l
          ∀x:int. is_in(x, nil) = false
          ∀x, y:int. ∀l:int list.
                                is_in(x, y::l) = true ⇔
                                (x = y ∨ is_in(x, l) = true)

SORT1 =
  enrich INTLIST by
  opns    is_sorted : int list → bool
          sort : int list → int list
  axioms  is_sorted(nil) = true
          ∀x:int. ∀l:int list.
                                is_sorted(x::l) = true ⇔
                                ((∀y:int. is_in(y, l) = true ⇒ po(x, y) = true)
                                ∧ is_sorted(l) = true)
          ∀l:int list. is_sorted(sort(l)) = true
          ∀l:int list. ∀x:int. is_in(x, l) = is_in(x, sort(l))

SORT = hide opns is_sorted in SORT1

```

We will work with this version of these specifications throughout the rest of the paper.

In examples throughout the rest of the paper, as in the above specifications, all signatures are taken to implicitly include all of the built-in type and value names of SML and all algebras extend the interpretation of those names given by the SML semantics. \square

3. Semantics of Specifications

Any specification language must be given a precise, formal semantics. The very concept of “correct” program is meaningless in the absence of a definition of what it is supposed to compute, and a specification can only provide such a definition if it has an unambiguously-defined meaning.

Before we start assigning meanings to specifications, it is necessary to decide what kind of mathematical objects to use to represent the meanings of specifications, i.e. to decide what specifications *denote*. Whatever the full answer is, a specification at least determines the underlying signature of the specified program. For any specification SP , we write this signature as $Sig(SP)$. Then, one may attempt to give a semantics of specifications on (at least) three different levels:

- Presentation level: a specification SP denotes a set of sentences over $Sig(SP)$ (this set may be required to be finite or at least recursive or recursively enumerable). At this level, the meaning of a specification is close to the syntactic form in which specifications are written; the semantics extracts the axioms, resolves references to other specifications, etc.
- Theory level: a specification SP denotes a *theory* over $Sig(SP)$, that is, a set of $Sig(SP)$ -sentences that is closed under logical consequence. This theory is much larger¹ than the set of axioms that are explicitly given in the specification. It is always infinite¹, usually not recursive and sometimes not recursively enumerable; thus the meaning of a specification is no longer strictly syntactic. The semantics performs the closure under logical consequence.
- Model-class level: a specification SP denotes a class of $Sig(SP)$ -algebras. At this level, the meaning of a specification is entirely non-syntactic, except for the signature part. The semantics abstracts away from the axioms, taking into account only their possible realizations.

The ultimate role of any specification is to describe a class of programs which we want to view as its correct realizations. Since we have already decided to model programs as algebras, *specifications ultimately determine classes of algebras*. Given the natural mappings from presentations to theories and from theories to model classes, this holds whichever one of these three levels we choose for the semantic domain.

For any specification SP , the semantics of SP determines the class of all *models* of SP , denoted by $\llbracket SP \rrbracket \subseteq Alg(Sig(SP))$ ². This class contains algebras that model programs which are considered to be correct realizations of SP . (There is a subtle issue involved in ensuring that *all* such algebras are admitted as models; see Section 9.) This semantics determines a notion of logical consequence of a specification: a specification SP entails a sentence φ , written $SP \models \varphi$, if φ holds in every model of SP .

Of course, a specification SP may admit a number of different program behaviours, and hence we cover so-called *loose* specifications. Or it might admit no models at all, in which case it is called *inconsistent*.

¹ Of course, this depends on the logic involved, but for example in equational logic every theory contains all the trivially true sentences like $\forall x:s. x = x$.

² Note the overloading of the semantic brackets: for a program P , $\llbracket P \rrbracket$ is an algebra, while for a specification SP , $\llbracket SP \rrbracket$ is a class of algebras.

The semantics of a specification formalism is usually presented by giving a number of semantic clauses, one for each specification-building operation. Each clause defines the meaning of a specification built using the given operation in terms of the meanings of its component specifications. This style of presentation gives a compositional semantics.

The following defines the models of specifications formed using the three specification-building operations that were informally presented earlier; cf. e.g. [SaT88a]. We omit the obvious context conditions which require that Σ as defined in each case is a well-formed signature.

Basic specifications:

$$\llbracket \text{sorts } S \text{ opns } \Omega \text{ axioms } \Phi \rrbracket = \{A \in \text{Alg}(\Sigma) \mid A \models \Phi\}$$

where $\Sigma = \text{Sig}(\text{sorts } S \text{ opns } \Omega \text{ axioms } \Phi)$ is the signature having sorts S and operations Ω .

Enrichment:

$$\begin{aligned} \llbracket \text{enrich } SP \text{ by sorts } S \text{ opns } \Omega \text{ axioms } \Phi \rrbracket = \\ \{A \in \text{Alg}(\Sigma) \mid A|_{\text{Sig}(SP)} \in \llbracket SP \rrbracket \text{ and } A \models \Phi\} \end{aligned}$$

where $\Sigma = \text{Sig}(\text{enrich } SP \text{ by sorts } S \text{ opns } \Omega \text{ axioms } \Phi)$ is the signature $\text{Sig}(SP)$ with additional sorts S and operations Ω .

Hiding:

$$\llbracket \text{hide sorts } S \text{ opns } \Omega \text{ in } SP \rrbracket = \{A|_{\Sigma} \mid A \in \llbracket SP \rrbracket\}$$

where $\Sigma = \text{Sig}(\text{hide sorts } S \text{ opns } \Omega \text{ in } SP)$ is the signature $\text{Sig}(SP)$ with sorts S and operations Ω removed.

Since specifications denote classes of algebras, specification-building operations semantically correspond to functions mapping classes of algebras to classes of algebras. Each of the definitions above amounts to the definition of such a function (a nullary one, in the case of basic specifications).

Example 1 (continued). The semantics of the specification-building operations can be used to calculate the meanings of specifications like those in Example 1 (Section 2), with the proviso given there concerning the built-in types and values of SML. For example, $\text{Sig}(\text{INTLIST})$ extends the built-in type and value names of SML by *po*, *head*, *tail* and *is_in*, and $\llbracket \text{INTLIST} \rrbracket$ is the class of algebras over this signature that extend the interpretation of the built-in names given by SML with operations *po*, *head*, *tail* and *is_in* defined in such a way that the axioms of *INTLIST* are satisfied. Then, $\text{Sig}(\text{SORT1})$ extends $\text{Sig}(\text{INTLIST})$ by *is_sorted* and *sort*, and $\llbracket \text{SORT1} \rrbracket$ is the class of algebras over this signature that enrich the algebras in $\llbracket \text{INTLIST} \rrbracket$ so that the axioms of *SORT1* are satisfied. Finally, $\text{Sig}(\text{SORT})$ extends $\text{Sig}(\text{INTLIST})$ by *sort* only, and $\llbracket \text{SORT} \rrbracket$ is the class of algebras over this signature that result from the algebras in $\llbracket \text{SORT1} \rrbracket$ by removing the interpretation of *is_sorted*. \square

3.1. Model Classes vs Theories

For any signature Σ , there is a Galois connection between classes of Σ -algebras and sets of Σ -sentences, assigning to any set of sentences the class of all algebras that satisfy them, and to any class of algebras the set of all sentences that

hold in them (see [GoB92]). The “closed” elements of this Galois connection are theories; these are in one-to-one correspondence with closed (i.e., definable by sets of sentences) classes of algebras. It follows from this that the theory level is less expressive as a semantic domain for specifications than either the presentation or the model-class level. The latter two are, however, incomparable: there are properties that can be naturally studied at the presentation level (for example, finiteness of an axiomatization) with no natural counterpart at the model-class level, and vice versa.

It is not immediately obvious that working at the model-class level brings any essential benefits over working with closed classes of algebras only, or equivalently, working at the theory level. It is not clear whether non-closed classes of algebras ever arise as meanings of specifications; even if they do arise, it is not clear whether this makes any difference for the use of specifications. The following example, built in the institution of equational logic, exhibits both of these phenomena:

$$SP \left\{ \begin{array}{l} \text{enrich} \\ SP_1 \left\{ \begin{array}{l} \text{hide opns } a \text{ in} \\ SP_0 \left\{ \begin{array}{l} \text{sorts } s, s' \\ \text{opns } a : s \\ b, c : s' \end{array} \right. \\ \text{by axioms } \forall x:s. b = c \end{array} \right. \end{array} \right.$$

This example relies on the following well-known fact [GoM85]: $\forall x:s. b = c$ does not imply $b = c$, although it implies $b = c$ for $\text{Sig}(SP)$ -algebras with non-empty carrier of sort s .

Now, according to the above definitions, $\llbracket SP_0 \rrbracket$ is the class of all algebras (over the indicated signature) and $\llbracket SP_1 \rrbracket$ consists of all algebras that are reducts of $\text{Sig}(SP_0)$ -algebras, obtained by removing the operation name a (but of course not its value). Consequently, $\llbracket SP_1 \rrbracket$ contains only those algebras having a non-empty carrier of sort s . Then, selecting from $\llbracket SP_1 \rrbracket$ the algebras that satisfy $\forall x:s. b = c$ yields the class $\llbracket SP \rrbracket$ — and all these algebras satisfy $b = c$ (since for the algebras in $\llbracket SP_1 \rrbracket$, $b = c$ follows from $\forall x:s. b = c$). Thus, under the model-class interpretation, the property $b = c$ is a consequence of the specification SP .

On the other hand, at the theory level, SP_0 would clearly have to denote the trivial equational theory containing only the equational tautologies, and so would SP_1 (there are no equations capable of expressing the fact that a carrier is non-empty). Then, the additional axiom $\forall x:s. b = c$ in the context of the theory denoted by SP_1 does not entail the equation $b = c$. Thus, under the theory-level interpretation, $b = c$ is *not* a consequence of the specification SP .

This discrepancy (and similar examples one may construct without relying on the “empty carriers” phenomenon) faces us with the necessity to choose between theories and classes of algebras as the basic semantic domain for specifications. The choice is obvious: the objects of ultimate interest here are programs, which are modelled as algebras, while axioms and theories are nothing more than logical means for describing them. In our view, the lack of agreement between theories and classes of algebras clearly demonstrates that theories are *not* in general adequate as denotations of specifications. But see [DGS93] for a different point of view.

The alert reader may have noticed that the above example depends crucially on the use of equational theories. If we reinterpret the example in the institution

of first-order predicate logic with equality, then the class $\llbracket SP_1 \rrbracket$ becomes definable (by the sentence $\exists x:s.true$) and the discrepancy between the theory level and the model class level semantics of SP disappears. This is an instance of a general phenomenon: as the expressive power of the logical system in use increases, the gap between the theory level and the model class level semantics narrows. For example, in the institution of second-order logic [HoS96], any class of models of a specification built using the specification-building operations presented above is definable by a set of axioms. The presence or absence of such a gap also depends on the expressive power of the specification-building operations in use. For example, if we use only basic specifications and **enrich** (leaving out hiding, as in Larch [GuH93] and ACT ONE [EhM85]), then there is no gap, no matter what institution we use.

4. Specification Engineering

The point of constructing a specification is so that it may be used to define a programming task by precisely delimiting the range of program behaviours that are to be regarded as permissible. The initial formal specification of requirements of a system thereby provides a reference point with respect to which all subsequent development activity is conducted. Specifications of system components play a similar role, but also serve to mediate proofs of correctness of systems containing them: a system (or sub-system) is proved to correctly implement its specification on the basis of those properties of components on which it depends that are recorded in their specifications. For these reasons, in the rest of this paper a formal specification of requirements is regarded as the starting point of system development.

There are serious problems involved in beginning with a formal specification of requirements, its desirability notwithstanding. Perhaps the most obvious problem is how to obtain a formal specification which accurately reflects the needs of the client. A program that is correct with respect to an incomplete or inaccurate specification of requirements is not of much use! This issue will be addressed in the remainder of this section. Another problem is that in real life, the requirements that any moderately complex system are expected to fulfill are subject to continual change. It follows that any fixed specification, formal or informal, can at best reflect a “snapshot” of the client’s needs. This suggests that the picture we present here needs to be augmented to accommodate changes in requirements, and that mechanisms are required to ensure that code (and proofs of correctness) keep in step with changes in requirements, cf. [GoL95].

The problem of writing the original requirements specification and ensuring that it is an accurate reflection of needs is the topic of “requirements engineering” [Dav90]. For some work on formal requirements analysis, see [RAH94, Li94, AsR95]. As suggested in the previous section, a key factor in facilitating the production of formal specifications is the provision of well-designed specification languages with good structuring operations allowing specifications to be built and understood in a systematic, modular fashion. Once a formal specification is obtained, the problem of checking that it is “correct” is of a different character from the problems treated in the remainder of this paper. Given a formal specification, it is possible, at least in principle, to prove (or disprove) that an alleged realization correctly implements it; this process is called *verification*. In contrast, an initial formal specification of requirements can at best be checked

for conformance with an informal written specification. Sometimes there will be no written specification at all and the formal specification can only be checked against the unwritten intentions of the client. The term *validation* is used to refer to the process of evaluating a specification against the client's written or unwritten informal requirements.

Just because a formal specification is precise and unambiguous does not mean that it is more likely than an informal specification to reflect needs accurately. Indeed, experience shows that problems uncovered by validation are often due to bugs in the formal specification rather than to errors in the informal specification. On the other hand, the process of writing a formal specification normally uncovers gaps or ambiguities in the informal specification. This means that validation is not merely a matter of checking that the formal specification accurately records what is already present in the informal specification; it is an iterative process which involves adjustment of both formal and informal specifications, and sometimes checking with the client to clarify needs. Since the cost of resolving problems with the requirements specification late in the development lifecycle may be extremely high, the production of a formal specification of requirements is regarded as a cost-effective activity, even if the resulting formal specification is not used in later stages of development [Som92].

One way of increasing confidence that a formal specification expresses what is required is to enable the client to “play” with it, in order to test whether or not the specification indeed expresses the properties he expects. A traditional approach to this is to engage a team of programmers to build a *prototype*, a quickly assembled but necessarily bad and simplified realization. This can then be given to the client to test. Of course, such an approach is indispensable for some aspects of the software to be developed. For example, there can be no better way to test a user interface than by playing with some version of it; going through sample sessions with such a system seems to be the only way for a user to get a feel for what working with the system will be like. In general, however, the prototyping approach has a number of disadvantages. First, it involves some extra work to produce a system that is then thrown away. More importantly, if the original specification is loose (and it usually is) then any prototype will incorporate choices between the alternative behaviours permitted by the specification, and these choices need not necessarily be mirrored in the final implementation. Consequently, the user may conclude that the system will have some properties that are not ensured by the specification at all, and this undermines the sense of the whole exercise. See [HaJ89] for further convincing arguments in this direction.

The overhead of prototyping may be avoided through the use of a rapid prototyping system like RAP [Hu85]. This demands that requirements specifications be written in an executable specification language, not far from high-level programming languages like Standard ML [Pau91]. In the fundamental trade-off between executability and expressiveness, it is clearly the latter that is of central importance in a language intended for writing requirements specifications, so such a strong restriction seems highly undesirable.

We believe that for many purposes prototyping should be replaced by *theorem proving* (see [GuH80] for a similar observation). To check whether a given specification indeed embodies a desirable property, it seems most appropriate to state this property explicitly and then try to prove that it is a consequence of the specification. This is the most general form of specification testing; the more usual approaches via rapid prototyping, symbolic evaluation, term rewriting etc.

can easily be seen as special cases, or rather as special techniques of theorem proving applicable in particular situations.

Example 1 (continued). In *INTLIST*, the axioms for *head*, *tail* and *is_in* virtually constitute a prototype implementation in e.g. SML. In a prototype, we would be able to evaluate expressions like *head(tail([2,4]))*, *head(tail([2,4,2]))*, *head(tail([2,4,5,8]))*, obtaining an integer value (4 in all these cases). However, rather than testing all these instances, we are able to *prove* directly from the specification a more general fact:

$$INTLIST \models \forall l:\text{int list}. \text{head}(\text{tail}(2::4::l)) = 4$$

Please note that some rapid prototyping systems allow the user to do somewhat more than evaluating just ground instances of *head(tail(2::4::l))*. For example, in RAP [Hu85], we could in fact evaluate *head(tail(2::4::l))* obtaining 4, as expected.

We can also prove the following fact:

$$SORT \models \text{po}(1,2) = \text{true} \implies \text{head}(\text{sort}([1,2,1])) = 1$$

However:

$$\begin{aligned} SORT &\not\models \text{po}(1,2) = \text{true} \implies \text{sort}([1,2,1]) = [1,1,2] \\ SORT &\not\models \text{po}(1,2) = \text{true} \implies \text{sort}([1,2,1]) = [1,2,2] \end{aligned}$$

even though a naive prototype implementation would probably satisfy one of these two equations. This would be misleading and potentially dangerous since *SORT* is loose: it does not specify whether or not *sort* should preserve repetitions. Sorting functions yielding either of these two results would be acceptable, and so would *sort*([1,2,1]) = [1,2,2]. \square

5. Proof Systems for Specifications

The above discussion indicated a need for formal proof systems for deriving consequences of specifications. Proof is also required for verifying the correctness of refinement steps, see below. There are two levels at which proof is necessary: first, we have to be able to derive consequences of sentences in the underlying institution ($\Phi \vdash \varphi$); second, we have to be able to derive consequences of a specification built in a structured way ($SP \vdash \varphi$). The first problem is familiar from logic, but the second has received much less attention. Here are inference rules that allow such consequences to be derived from specifications built using the specification-building operations introduced above; see [SaB83, SaT88a, Far92, Wir93].

$$\begin{array}{c} \frac{\varphi \in \Phi}{\text{sorts } S \text{ opns } \Omega \text{ axioms } \Phi \vdash \varphi} \\[10pt] \frac{SP \vdash \varphi}{\text{enrich } SP \text{ by sorts } S \text{ opns } \Omega \text{ axioms } \Phi \vdash \varphi} \\[10pt] \frac{\varphi \in \Phi}{\text{enrich } SP \text{ by sorts } S \text{ opns } \Omega \text{ axioms } \Phi \vdash \varphi} \end{array}$$

$$\begin{array}{c}
\frac{SP \vdash \varphi}{\text{hide sorts } S \text{ opns } \Omega \text{ in } SP \vdash \varphi} \quad \varphi \text{ is a } (\text{Sig}(SP) \setminus \langle S, \Omega \rangle)\text{-sentence} \\
\\
\frac{SP \vdash \varphi_1 \quad \cdots \quad SP \vdash \varphi_n \quad \{\varphi_1, \dots, \varphi_n\} \vdash \varphi}{SP \vdash \varphi}
\end{array}$$

The last of these rules constitutes the link between the two levels of proof.

Example 1 (continued). Here is an example of how these rules may be used in the proof of a simple consequence of the specification *SORT*. For simplicity, all universal quantifiers are omitted, and we tacitly α -convert where necessary to avoid variable clashes.

$$\frac{\begin{array}{ccc} (1) & (2) & (3) \\ \hline \text{SORT1} \vdash \text{is_in}(2, \text{sort}(0::2::l)) = \text{true} & & \\ \hline \text{SORT} \vdash \text{is_in}(2, \text{sort}(0::2::l)) = \text{true} \end{array}}{}$$

where (1) is the following derivation:

$$\frac{\begin{array}{ccc} (1.1) & & (1.2) \\ \hline \text{INTLIST} \vdash \text{is_in}(2, 0::2::l) = \text{true} & & \\ \hline \text{SORT1} \vdash \text{is_in}(2, 0::2::l) = \text{true} \end{array}}{}$$

with (1.1) being:

$$\frac{\begin{array}{c} (\text{is_in}(x, y::l') = \text{true} \iff (x = y \vee \text{is_in}(x, l') = \text{true})) \\ \in \text{axioms of INTLIST} \end{array}}{\text{INTLIST} \vdash \text{is_in}(x, y::l') = \text{true} \iff (x = y \vee \text{is_in}(x, l') = \text{true})}$$

and (1.2):

$$\frac{\{\text{is_in}(x, y::l') = \text{true} \iff (x = y \vee \text{is_in}(x, l') = \text{true})\} \vdash}{\text{is_in}(2, 0::2::l) = \text{true},}$$

(2) is:

$$\frac{\begin{array}{ccc} (2.1) & & (2.2) \\ \hline \text{SORT1} \vdash \text{is_in}(2, 0::2::l) = \text{is_in}(2, \text{sort}(0::2::l)) & & \end{array}}{}$$

with (2.1) being:

$$\frac{(\text{is_in}(x, l') = \text{is_in}(x, \text{sort}(l'))) \in \text{axioms of SORT1}}{\text{SORT1} \vdash \text{is_in}(x, l') = \text{is_in}(x, \text{sort}(l'))}$$

and (2.2):

$$\{\text{is_in}(x, l') = \text{is_in}(x, \text{sort}(l'))\} \vdash \text{is_in}(2, 0::2::l) = \text{is_in}(2, \text{sort}(0::2::l)),$$

and (3) is the following entailment:

$$\{\text{is_in}(2, 0::2::l) = \text{true}, \text{is_in}(2, 0::2::l) = \text{is_in}(2, \text{sort}(0::2::l))\} \vdash \text{is_in}(2, \text{sort}(0::2::l)) = \text{true}.$$

□

The above rules provide a *sound* extension to any proof system for the underlying institution: if $\Phi \models \varphi$ whenever $\Phi \vdash \varphi$ for all Φ and φ , then $SP \models \varphi$ whenever $SP \vdash \varphi$ for all SP and φ . *Completeness* ($SP \models \varphi$ implies $SP \vdash \varphi$) is harder to achieve. Even if the proof system for the underlying institution is complete, the above rules do not in general yield a complete proof system for

consequences of structured specifications (but a complete system *is* obtained for institutions satisfying an appropriate interpolation property [Cen94]). Whenever there is a discrepancy between model class level and theory level semantics as discussed in Section 3.1, no complete compositional proof system may be given. This does not exclude the existence of non-compositional complete proof systems that “massage” the structure of specifications in the course of proof, see [Far92, Wir93].

Good theorem provers that implement such proof systems are needed. In addition to proof search procedures used in work on automatic theorem proving, they should include heuristics that exploit the structure of specifications to guide proof search, see [SaB83, HST94].

Example 1 (continued). A theorem prover attempting to prove that

$$SORT \vdash \forall l:\text{int list. is_in}(2, 0::2::l) = \text{true}$$

should not waste time searching through the consequences of the axioms added in *SORT1*, but should go straight to the level of *INTLIST* where most of the work of the proof needs to be done. The following heuristics would provide the necessary guidance:

- To prove **hide** ... **in** $SP \vdash \varphi$, try to prove $SP \vdash \varphi$.
- To prove **enrich** SP **by** **sorts** S **opns** Ω **axioms** ... $\vdash \varphi$, if φ doesn't use any of the new operations in Ω , try to prove $SP \vdash \varphi$.

The latter of the above simple heuristics, even though it does not yield a complete proof method, in practice often helps to greatly reduce the proof search space (cf. [SaB83, HST94]). \square

It would also be extremely useful for a theorem prover, in the case where it fails to find a proof, to provide the user with readable information on where the proof attempt breaks down (see e.g. LP [GuH93]), and perhaps even how the specification may be augmented to make the proof go through — a desirable feature which few contemporary theorem provers exhibit.

6. Program Development

Given a specification SP , the programming task it defines is to construct a program P that is a correct realization of SP , that is, such that $\llbracket P \rrbracket \in \llbracket SP \rrbracket$.

There can be no universal recipe that would ensure successful development of a program implementing a given specification. All we can hope to offer are methodologies, and particular techniques and heuristics oriented towards specific problem areas.

Perhaps the most fundamental point is that it is neither easy nor desirable to leap in a single bound over the gap between a high-level user-oriented requirements specification and the realm of programs full of technical decisions and algorithmic details. An attractive alternative is to proceed systematically in a stepwise fashion, gradually enriching the original requirements specification with more and more detail, incorporating more and more design and implementation decisions. Such decisions include choosing between the options of behaviour left open by the specification, between the algorithms that realize this behaviour, between data representation schemes, etc. Each such decision is recorded separately,

as a separate step hopefully consisting of a local modification to the specification. Developing a program from a specification then proceeds via a sequence of such small, easy to understand and easy to verify steps:

$$SP_0 \rightsquigarrow SP_1 \rightsquigarrow \cdots \rightsquigarrow SP_n$$

In such a chain, SP_0 is the original requirements specification and $SP_{i-1} \rightsquigarrow SP_i$ for any $i = 1, \dots, n$ is an individual *refinement step*. The aim is to reach a specification (here, SP_n) that is an exact description of a program in full detail, with all the technical decisions incorporated (it may simply *be* a program, if our specification formalism is rich enough).

Example 1 (continued). The following adds to the specification *SORT* the decision that the sorting operation *sort* should preserve the number of occurrences of elements so that the result is a permutation of the argument list:

```

SORTperm =
  hide opns count in
    enrich SORT by
      opns    count : int × int list → int
      axioms  ∀x:int. count(x, nil) = 0
              ∀x,y:int. ∀l:int list.
                  x ≠ y ⇒ count(x, y::l) = count(x, l)
              ∀x:int. ∀l:int list. count(x, x::l) = 1+count(x, l)
              ∀x:int. ∀l:int list. count(x, l) = count(x, sort(l))

```

Then we choose the algorithm (insertion sort) and “code” *sort* but, for illustrative purposes, we refrain at this stage from giving the “code” for the additional operation *insert* and leave it specified only.

```

INS = enrich INTLIST by
  opns    insert : int × int list → int list
  axioms  ∀x:int. ∀l:int list. ∃l1,l2:int list.
              insert(x, l) = l1@(x::l2) ∧ l = l1@l2
              ∧ (∀l1':int list. ∀y:int. l1 = l1'@[y] ⇒
                  po(y, x) = true)
              ∧ (∀l2':int list. ∀y:int. l2 = y::l2' ⇒
                  po(x, y) = true)

```

```

SORTins =
  hide opns insert in
    enrich INS by
      opns    sort : int list → int list
      axioms  sort(nil) = nil
              ∀x:int. ∀l:int list. sort(x::l) = insert(x, sort(l))

```

Finally, we “code” *insert*, preserving the “code” for *sort* :

```

INSdone =
  enrich INTLIST by
    opns    insert : int × int list → int list
    axioms  ∀x:int. insert(x, nil) = [x]
             ∀x,y:int. ∀l:int list.
                 po(x, y) = true ⇒
                     insert(x, y::l) = x::y::l
             ∀x,y:int. ∀l:int list.
                 po(x, y) = false ⇒
                     insert(x, y::l) = y::insert(x, l)

SORTdone =
  hide opns insert in
    enrich INSdone by
      opns    sort : int list → int list
      axioms  sort(nil) = nil
               ∀x:int. ∀l:int list. sort(x::l) = insert(x, sort(l))

```

The above constitutes a sequence of development steps:

$$SORT \rightsquigarrow SORT_{perm} \rightsquigarrow SORT_{ins} \rightsquigarrow SORT_{done}$$

SORTdone may be viewed as a final implementation of the original specification since the axioms in *INTLIST*, *INSdone* and *SORTdone* amount to SML code (this disregards the fact that *po* is only specified as a partial order, rather than being coded as a specific order relation). We will make this more explicit in the next section. \square

A formal definition of such refinement steps $SP \rightsquigarrow SP'$ must incorporate the requirement that any correct final realization of SP' must be a correct realization of SP . This leads to the following straightforward definition [SaW83, SaT88b]:

$$SP \rightsquigarrow SP' \quad \text{iff} \quad \llbracket SP' \rrbracket \subseteq \llbracket SP \rrbracket$$

(This presupposes that $\text{Sig}(SP) = \text{Sig}(SP')$.)

Example 1 (continued). The refinement steps in the above example satisfy the definition. This is trivial for the step $SORT \rightsquigarrow SORT_{perm}$, since *SORTperm* just adds a constraint on the class of models of *SORT*. For $SORT_{perm} \rightsquigarrow SORT_{ins}$, it is necessary to prove that to each model of *SORTins*, we can add *count* and *is_sorted* so that the axioms of *SORTperm* are satisfied. Since *count* and *is_sorted* are determined by the corresponding axioms in *SORTperm*, this amounts to proving that the “code” for *sort* entails the axioms of *SORTperm*, assuming that *insert* satisfies the axioms in *INS* and that *count* and *is_sorted* satisfy their axioms. Finally, $SORT_{ins} \rightsquigarrow SORT_{done}$ requires a proof that the “code” for *insert* in *INSdone* entails the axiom in *INS*. The reader is encouraged to check the details.

It is perhaps worth noticing that $\llbracket SORT_{perm} \rrbracket = \llbracket SORT_{ins} \rrbracket = \llbracket SORT_{done} \rrbracket$ (even though $\llbracket INS \rrbracket \neq \llbracket INS_{done} \rrbracket$, and *count*, hidden in *SORTperm*, is not even mentioned in *SORTins* and *SORTdone*); this means that the last two refinement steps are semantically trivial although this does not mean that the proofs are trivial. The reader may be worried by the fact that it then follows that, for example, $SORT_{done} \rightsquigarrow SORT_{perm}$. The notion of refinement is not fine enough to capture the sense in which *SORTdone* is “closer” to a program than *SORTperm* is. A

more elaborate notion of refinement, which provides a place to record “progress” towards a program, will be presented in the next section. \square

The definition of refinement ensures that the correctness of the final outcome of stepwise development may be inferred from the correctness of the individual refinement steps:

$$\frac{SP_0 \rightsquigarrow SP_1 \rightsquigarrow \cdots \rightsquigarrow SP_n \quad A \in \llbracket SP_n \rrbracket}{A \in \llbracket SP_0 \rrbracket}$$

The proof is by an easy induction on the length of the refinement sequence.

Notice that if the final specification SP_n represents an individual program P , i.e. $\llbracket SP \rrbracket = \{\llbracket P \rrbracket\}$, then the conclusion that $A \in \llbracket SP_0 \rrbracket$ for all $A \in \llbracket SP_n \rrbracket$ is just our original statement of the program development task: $\llbracket P \rrbracket \in \llbracket SP_0 \rrbracket$.

An indirect way to prove the correctness of the final outcome is to notice a stronger fact, namely that consecutive refinements can be composed (referred to as “vertical composability” [GoB80]):

$$\frac{SP \rightsquigarrow SP' \quad SP' \rightsquigarrow SP''}{SP \rightsquigarrow SP''}$$

The above gives a formal view of the stepwise development methodology. As mentioned before, there can be no universal recipe for coming up with useful refinements of a given specification — necessarily, this is the place where the developer’s invention is required. Once a refinement step is proposed, though, we should be able to prove it correct, that is, we should have some formalism for proving the inclusion between the corresponding model classes. Composing the proofs of all the steps involved in the development of a program from a specification gives a proof that the program is correct with respect to the original specification. But there seems to be no benefit in actually producing this proof: individual proofs of correctness of the individual steps are easier to produce and easier to understand than a single monolithic proof of correctness of the resulting program.

A formalism for proving correctness of refinement steps must of course incorporate a theorem prover for the underlying logic, and for proving consequences of structured specifications, as discussed above. A new need that arises here is that of proving entailments between two structured specifications (we write $SP' \models SP$ to state that every model of SP' is a model of SP , yet another formulation of $SP \rightsquigarrow SP'$ that we will use in this context). If the structures of SP and SP' match exactly (and the specification-building operations used are monotonic w.r.t. inclusion of model classes — this holds for all the specification-building operations in this paper and is typically the case for those considered elsewhere in the literature) then this problem may be reduced to proving that individual axioms (from SP) are consequences of certain specifications (parts of SP') via the following fact, which is referred to as “horizontal composability” [GoB80] for the specification-building operation op :

$$\frac{SP_1 \rightsquigarrow SP'_1 \quad \cdots \quad SP_n \rightsquigarrow SP'_n}{op(SP_1, \dots, SP_n) \rightsquigarrow op(SP'_1, \dots, SP'_n)}$$

Unfortunately, the structures of the two specifications need not coincide, which makes such a reduction very non-trivial. The only work on this important problem we are aware of is [Far92, Wir93].

Example 1 (continued). The refinements in Example 1 illustrate the point that

the structure of the final implementation differs from that of the original specification, even though the difference is only that different auxiliary operations are used (*is_sorted* versus *insert*). The essential change happens in the step $SORTperm \rightsquigarrow SORTins$. There is no way to build this refinement using the horizontal composability rule: stripping off the hiding operations from both $SORTperm$ and $SORTins$ (naively disregarding the fact that different things are hidden) yields two incomparable specifications. See [Far92, Wir93] for proof rules that allow the user to handle this situation in a different, non-compositional way. \square

Horizontal composability should not be misread as a directive to decompose the task of realizing a specification $SP = op(SP_1, \dots, SP_n)$ into separate tasks to realize each of SP_1, \dots, SP_n . It is possible for the design decisions taken in the solutions of these separate tasks to conflict so that even once we have obtained realizations of SP_1, \dots, SP_n , it might not be possible to combine these to form a realization of SP .

Example 2. Consider the following specification:

$$SPc = \begin{array}{ll} \text{enrich opns} & c : \text{int} \\ & \text{axioms } 1 < c < 15 \\ \text{by axioms} & 10 < c < 27 \end{array}$$

Since

$$\left(\begin{array}{ll} \text{opns} & c : \text{int} \\ \text{axioms} & 1 < c < 15 \end{array} \right) \rightsquigarrow \left(\begin{array}{ll} \text{opns} & c : \text{int} \\ \text{axioms} & 1 < c < 12 \end{array} \right)$$

and

$$\left(\begin{array}{ll} \text{opns} & c : \text{int} \\ \text{axioms} & 10 < c < 27 \end{array} \right) \rightsquigarrow \left(\begin{array}{ll} \text{opns} & c : \text{int} \\ \text{axioms} & 14 < c < 20 \end{array} \right)$$

we also have

$$SPc \rightsquigarrow \left(\begin{array}{ll} \text{enrich opns} & c : \text{int} \\ & \text{axioms } 1 < c < 12 \\ \text{by axioms} & 14 < c < 20 \end{array} \right)$$

However, even though SPc is consistent, and both of the resulting component specifications are consistent as well, the resulting composed specification to which SPc is refined is inconsistent!

This happens because the two specification arguments to the **enrich** operation implicitly share a loosely specified part ($c : \text{int}$ in the example). If the decisions constraining this common part in separate developments of the two specifications are different, as above, then putting the resulting specifications together may yield inconsistency. This is of course a contrived example but the same phenomenon arises in more realistic situations. \square

An issue which may seem worrying here is that we have not put into our definition of refinement any requirement that the refined specification is consistent. Indeed, this can be seen as a problem, since an inconsistent specification cannot be implemented by any program, and so it opens a blind alley in the program development process. From this point of view, it would be worthwhile to be able to check consistency of each specification as soon as it is formulated. Unfortunately, in general (for any sufficiently powerful specification framework)

this is an undecidable property. Fortunately, inconsistency of specifications cannot lead to incorrect programs: if we arrive at a program at some point in the development process, then this program is by definition consistent (it has a unique model) and consequently, all the specifications leading to it must have been consistent as well.

The proposed methodology of stepwise refinement does not and cannot be expected to guarantee success. Apart from inconsistencies, there are many sources of blind alleys and failures in the development process: there might be no computable realization of a specification, there might be no “computationally feasible” realization, we might not be clever enough to find a realization, we might run out of money to finish the project, etc.

Example 3. Consider a specification of natural numbers with a pre-ordering specified by the sentence:

$$m < n = \text{true} \iff \forall x:\text{nat}. M_m \downarrow x \implies M_n \downarrow x$$


where for all natural numbers k and x , the predicate $M_k \downarrow x$ is specified to mean that the Turing machine with Gödel number k terminates on input x . This specification is consistent but it has no computable models since the halting problem is undecidable. \square

The main feature of the methodology we really can ensure is its *safety*: if we arrive at a program, then it is a correct realization of the original specification.


Some refinement steps are more or less routine. For instance, there are standard ways of implementing many data abstractions (e.g. sets, queues) and standard ways of decomposing problems into simpler sub-problems (e.g. “divide and conquer”). Such refinement steps can sometimes be described schematically by means of so-called *transformation rules* such that any instance is guaranteed to be correct provided certain conditions are met. This reduces the burden of proving correctness of refinement steps: a proof that the transformation rule is correct is of course required, but this only needs to be done once for each rule. Then a simpler proof is required to show that the applicability conditions attached to the rule are satisfied, each time the rule is instantiated. The use of transformation rules also avoids the need for the programmer to come up with the idea for each refinement step from scratch. The need for creativity is not eliminated, of course: the application of transformation rules often requires invention of functions or conditions that do not appear in the specification being transformed to be substituted for schema variables. There has been a great deal of work on the transformational method of software development, much of it focussed on improving programs rather than on developing programs from specifications. A recent reference is [HKB93].

7. Constructor Implementations

The simple notion of specification refinement is mathematically elegant and powerful enough (in the context of a sufficiently rich specification language) to handle all concrete examples of interest. However, it is not very convenient to use in practice. During the process of developing a program, the successive specifications incorporate more and more details arising from successive design decisions. Thereby, some parts become fully determined, and remain unchanged as a part of the specification until the final program is obtained.



It is more convenient to avoid such clutter by separating the finished parts from the specification, putting them aside, and proceeding with the development of the unresolved parts only.



It is important for the finished parts $\kappa_1, \dots, \kappa_n$ to be independent of the particular choice of realization for what is left: they should act as constructions extending any realization of the unresolved part to a realization of what is being refined.

Example 1 (continued). An instance of the situation illustrated above may be found in the consecutive refinement steps $SORT_{perm} \rightsquigarrow SORT_{ins} \rightsquigarrow SORT_{done}$: the “code” for *sort* introduced in $SORT_{ins}$, and the operation that hides *insert*, are still present in the same form in $SORT_{done}$. \square

Each κ_i above amounts to what is known as a *parameterized program* [Gog84] with input interface SP_i and output interface SP_{i-1} . Given a program P that is a correct realization of SP_i , the parameterized program κ_i may be instantiated to yield a program $\kappa_i(P)$ that realizes SP_{i-1} . A programming language that supports stepwise development in the style suggested here needs to provide syntax and modularisation facilities for defining parameterized programs and their instantiations. For example, in the Standard ML programming language [Pau91] parameterized programs are called *functors*, and instantiation amounts to *functor application*.³ In Modula-3 [Nel91], parameterized programs are called *generic modules*. Once the development is finally finished (that is, when nothing is left unresolved, as above) we can successively instantiate the parameterized programs $\kappa_n, \dots, \kappa_1$ to obtain a correct realization of the original specification SP_0 .

Semantically, each parameterized program κ_i defines a function (which we will call a *constructor*)⁴ on algebras, $\llbracket \kappa_i \rrbracket : Alg(Sig(SP_i)) \rightarrow Alg(Sig(SP_{i-1}))$, and instantiation is simply function application: if $\llbracket P \rrbracket \in Alg(Sig(SP_i))$, then $\llbracket \kappa_i(P) \rrbracket = \llbracket \kappa_i \rrbracket(\llbracket P \rrbracket)$. In practice, κ_i provides a definition of the components (carriers and operations) of a $Sig(SP_{i-1})$ -algebra, given the components of a $Sig(SP_i)$ -algebra.

³ In the following we disregard the fact that functor application in SML is not guaranteed to terminate. The technicalities may be modified to capture this by modelling parameterized programs as *partial* (rather than total) functions and adding the obvious definedness condition in the definition of constructor implementation [SaT89]. We resist the temptation to adopt this slightly more complex approach for the sake of clarity of presentation.

⁴ Constructors should not be confused with *value constructors* like `nil` and `::` in SML and similar programming languages.

Example 1 (continued). Consider the refinement $SORT_{perm} \rightsquigarrow SORT_{ins}$ in which “code” for *sort* is first introduced. Using a notation like that of Standard ML, a parameterized program corresponding to this step can be expressed as follows:

```

functor K1(X:INS):SORTperm =
  struct
    open X
    fun sort(nil) = nil
      | sort(x::l) = insert(x,sort(l))
  end

```

(The effect of the declaration “open X” is to add the types and values in the parameter X to the context, allowing the use of names like *insert* and *head* in place of the qualified names *X.insert* and *X.head*. The reader is asked to find the obvious correspondence between the names used here and those used — in a different font — in the specifications.)

Recall that *INS* is the part of *SORT_{ins}* that remains after “peeling off” *sort* and the operation of hiding *insert*, the part of the specification whose implementation is fixed in this step. Notice that the functor definition provides not only code for *sort* but also (implicitly) realizes the hiding of *insert* since *insert* is not present in the functor result signature.

The next refinement step, $SORT_{ins} \rightsquigarrow SORT_{done}$, which introduces code for *insert*, corresponds to the following parameterized program:

```

functor K2(X:INTLIST):INS =
  struct
    open X
    fun insert(x,nil) = [x]
      | insert(x,y::l) = if po(x,y) then x::y::l
                          else y::insert(x,l)
  end

```

The code for *sort*, which in the original refinement step was still present in *SORT_{done}*, has been dealt with in the previous step. Thus in this step we are able to focus on what remains, namely the *insert* operation, without the distraction of the surrounding context.

The axioms in *INTLIST* may be translated directly into SML code, and we can choose a particular realization to implement *po*, giving the following parameterized program:

```

functor K3(X:EMPTY):INTLIST =
  struct
    fun po(n,m:int) = n <= m
    fun head(x::_) = x
    fun tail(_::l) = l
    fun is_in(_,nil) = false
      | is_in(x,y::l) = (x=y) orelse is_in(x,l)
  end

```

Here, *EMPTY* stands for the empty SML signature *sig end*.

To finish the example, we need to provide a parameterized program corresponding to the refinement step $SORT \rightsquigarrow SORT_{perm}$. Since all that is done in

this step is to impose a (non-constructive) restriction on the class of permissible realizations of *sort*, this is trivial:

functor $K0(X: \text{SORTperm}): \text{SORT} = X$

□

The above considerations motivate a more elaborate version of the notion of refinement of the previous section, known as *constructor implementation* [SaT88b]. We write $SP \rightsquigarrow_{\kappa} SP'$ to say that a specification SP' implements a specification SP via κ , where κ is a parameterized program denoting a constructor $\llbracket \kappa \rrbracket : \text{Alg}(\text{Sig}(SP')) \rightarrow \text{Alg}(\text{Sig}(SP))$, and define this as follows:

$$SP \rightsquigarrow_{\kappa} SP' \quad \text{iff} \quad \llbracket \kappa \rrbracket(\llbracket SP' \rrbracket) \subseteq \llbracket SP \rrbracket$$

Here, $\llbracket \kappa \rrbracket(\llbracket SP' \rrbracket)$ is the image of $\llbracket SP' \rrbracket$ under $\llbracket \kappa \rrbracket$.

Example 1 (continued). The following are examples of constructor implementations:

$$\begin{aligned} \text{SORT} &\rightsquigarrow_{K0} \text{SORTperm} \\ \text{SORTperm} &\rightsquigarrow_{K1} \text{INS} \\ \text{INS} &\rightsquigarrow_{K2} \text{INTLIST} \\ \text{INTLIST} &\rightsquigarrow_{K3} \text{EMPTY} \end{aligned}$$

The justification requires proofs similar to those sketched in Example 1 in Section 6 for the corresponding refinement steps. □

For each parameterized program κ we can (in principle) define a specification-building operation $\bar{\kappa}$ such that $\llbracket \bar{\kappa}(SP') \rrbracket = \llbracket \kappa \rrbracket(\llbracket SP' \rrbracket)$; then constructor implementations may be viewed as refinements ($SP \rightsquigarrow_{\kappa} SP'$ is just $SP \rightsquigarrow \bar{\kappa}(SP')$). Provided that we have means for reasoning about specifications built using these new operations, the correctness of constructor implementations may be established using proof techniques for refinements. Specifically, we need a way of deriving entailments of the form $\bar{\kappa}(SP') \models SP$; this boils down to proving properties of the components of programs built by κ .

The correctness of the final outcome of the stepwise development process may be inferred from the correctness of the individual constructor implementation steps:

$$\frac{SP_0 \rightsquigarrow_{\kappa_1} SP_1 \rightsquigarrow_{\kappa_2} \dots \rightsquigarrow_{\kappa_n} SP_n = \text{EMPTY}}{\llbracket \kappa_1(\kappa_2(\dots \kappa_n(\text{empty}) \dots)) \rrbracket \in \llbracket SP_0 \rrbracket}$$

where *EMPTY* is the empty specification over the empty signature and *empty* is its (empty) realization.

Example 1 (continued). In our example,

$$K0(K1(K2(K3(\text{empty}))))$$

yields a non-parameterized program (an SML structure) satisfying *SORT*. Here, *empty* stands for the empty SML structure `struct end`. □

Suppose that parameterized programs compose, that is, for any two parameterized programs κ and κ' such that the signature of the output interface of κ coincides with the signature of the input interface of κ' , there is a parameterized

program $\kappa;\kappa'$ with $\llbracket \kappa;\kappa' \rrbracket = \llbracket \kappa \rrbracket; \llbracket \kappa' \rrbracket$ (the latter semicolon stands for ordinary function composition, written in diagrammatic order). Then it is easy to see that constructor implementations (vertically) compose:

$$\frac{SP \rightsquigarrow_{\kappa} SP' \quad SP' \rightsquigarrow_{\kappa'} SP''}{SP \rightsquigarrow_{\kappa;\kappa'} SP''}$$

The requirement that parameterized programs can be instantiated is a weaker requirement than that parameterized programs be composable, even though any programming language with decent modularisation facilities should ensure the latter as well. In Standard ML, there is no explicit functor composition operation but the composite of two functors may easily be defined using functor application and abstraction.

As in the case of refinement, vertical composability is not necessary to ensure the correctness of the outcome of the development process. All we need is the condition inherent in the definition of constructor implementation, namely that implementations reflect realizations:

$$\frac{SP \rightsquigarrow_{\kappa} SP' \quad A' \in \llbracket SP' \rrbracket}{\llbracket \kappa \rrbracket(A') \in \llbracket SP \rrbracket}$$

Many approaches to implementation (see e.g. [EKM82, SaW82, Ore83]) make use of a restrictive kind of constructor defined by a parameterized program having a particular rigid form. Then the vertical composition of two implementations is required to yield an implementation of the same form, which is not always possible. The requirement that the composition of parameterized programs be forced into some given normal form corresponds to requiring programs to be written in a rather restricted programming language.

We have already mentioned that the internal structure of a requirements specification need not be mirrored by programs that realize it. This is why the definitions of refinement and constructor implementation above take no account of the structure of specifications. However, when developing a large program it is crucial to progressively decompose the job into smaller tasks that can be handled separately. Each task is defined by a specification, and solving a task means producing a program component that satisfies this specification. Once all tasks are solved, producing the final system is a simple matter of appropriately assembling these components.

A development step involving the decomposition of a programming task into separate subtasks is modelled using a constructor implementation with a multi-argument parameterized program (see [SST92]):

$$SP \rightsquigarrow_{\kappa} \langle SP_1, \dots, SP_n \rangle \quad \text{iff} \quad \llbracket \kappa \rrbracket(\llbracket SP_1 \rrbracket \times \dots \times \llbracket SP_n \rrbracket) \subseteq \llbracket SP \rrbracket$$

where $\llbracket \kappa \rrbracket : \text{Alg}(\text{Sig}(SP_1)) \times \dots \times \text{Alg}(\text{Sig}(SP_n)) \rightarrow \text{Alg}(\text{Sig}(SP))$ is an n -argument constructor (an n -argument function on algebras) describing a way to put models of SP_1, \dots, SP_n together to construct a model of SP (and, as before, we use the same notation $\llbracket \kappa \rrbracket$ to denote the corresponding image function). Now the development takes on a tree-like shape. The development is complete once a tree is obtained that has empty sequences (of specifications) as its leaves:

$$SP \xrightarrow{\kappa} \left\{ \begin{array}{l} SP_1 \xrightarrow{\kappa_1} \langle \rangle \\ \vdots \\ SP_n \xrightarrow{\kappa_n} \left\{ \begin{array}{l} SP_{n1} \xrightarrow{\kappa_{n1}} \{ SP_{n11} \xrightarrow{\kappa_{n11}} \langle \rangle \\ \dots \\ SP_{nm} \xrightarrow{\kappa_{nm}} \langle \rangle \end{array} \right\} \end{array} \right.$$

Then an appropriate instantiation of the parameterized programs in the tree yields a realization of the original requirements specification. The above development tree yields the program $\kappa(\kappa_1(), \dots, \kappa_n(\kappa_{n1}(\kappa_{n11}()), \dots, \kappa_{nm}()))$, with

$$\llbracket \kappa(\kappa_1(), \dots, \kappa_n(\kappa_{n1}(\kappa_{n11}()), \dots, \kappa_{nm}())) \rrbracket \in \llbracket SP \rrbracket$$

(We use an obvious notation $\kappa(P_1, \dots, P_n)$ for instantiation of n -ary parameterized programs, where $\llbracket \kappa(P_1, \dots, P_n) \rrbracket = \llbracket \kappa \rrbracket(\llbracket P_1 \rrbracket, \dots, \llbracket P_n \rrbracket)$.)

The structure of the final program is determined by the shape of the development tree, which is in turn determined by the decomposition steps. Each such step corresponds to what software engineers call a *design* specification (and what [GHW82] call an *organizational* specification): it defines the structure of the system by specifying its components and describing how they fit together. This style of development leads to modular programs, built from fully specified, correct and reusable components.

A complete development tree does not reflect the *process* of developing a system from a specification, which normally involves false starts, blind alleys and backtracking. It documents only the final outcome of this process, where all subtasks have been solved successfully. An incomplete development tree may be used to record a stage in the development process, so the development process corresponds to a sequence of such trees which culminates in a complete tree. Ideally, each tree in the sequence is an expansion of the previous one, but backtracking corresponds to deletion or alteration of parts of the tree that have already been filled in.

Example 1 (continued). We show a simple example of a decomposition using a modified version of the sorting specification above:

```

SORTonce =
  hide opns all_once in
    enrich SORT by
      opns    all_once : int list → bool
      axioms  all_once(nil) = true
               ∀x:int. ∀l:int list.
                 all_once(x::l) = true ⇔
                   all_once(l) = true ∧ is_in(x, l) = false
               ∀l:int list. all_once(sort(l)) = true

```

SORT_{once} just adds to **SORT** the requirement that the result of *sort* does not contain multiple occurrences of elements. Clearly, **SORT** \rightsquigarrow **SORT**_{once}.

Consider an additional specification that introduces a function specified to remove adjacent occurrences of the same element in a list:

```

NOSTUTTER =
  enrich INTLIST by
    opns    rem_stutter : int list → int list
    axioms   $\forall l, l1, l2 : \text{int list}. \forall x, y : \text{int}.$ 
                $\text{rem\_stutter}(l) = l1 @ (x :: y :: l2) \implies x \neq y$ 
                $\forall l : \text{int list}. \forall x : \text{int}.$ 
                $\text{is\_in}(x, l) = \text{is\_in}(x, \text{rem\_stutter}(l))$ 
                $\forall l, l1, l2 : \text{int list}.$ 
                $\text{rem\_stutter}(l) = l1 @ l2 \implies$ 
                    $\exists l3, l4 : \text{int list}.$ 
                    $l = l3 @ l4 \wedge \text{rem\_stutter}(l3) = l1$ 
                    $\wedge \text{rem\_stutter}(l4) = l2$ 

```

Now, the problem of implementing the specification *SORTonce* may be decomposed (perhaps not very efficiently, but certainly correctly) into the problems of implementing *SORTperm* and *NOSTUTTER*:

$$\text{SORTonce} \rightsquigarrow_{K4} \langle \text{SORTperm}, \text{NOSTUTTER} \rangle$$

where the parameterized program K4 is given as follows:

```

functor K4(X:SORTperm, Y:NOSTUTTER):SORTonce =
  struct
    open X
    fun sort(l) = Y.rem_stutter(X.sort(l))
  end

```

The specification *NOSTUTTER* can easily be implemented using the following functor:

```

functor K5(Z:INTLIST):NOSTUTTER =
  struct
    fun rem_stutter(nil) = nil
      | rem_stutter([x]) = [x]
      | rem_stutter(x::y::l) =
          if x = y then rem_stutter(y::l)
          else x::rem_stutter(y::l)
  end

```

Since we already have an implementation of *SORTperm* (obtained entirely independently from the development for *NOSTUTTER*) and of *INTLIST*, the development is complete and we can put all these together to obtain the following realization of the specification *SORTonce*:

$$K4(K1(K2(K3(\text{empty}))), K5(K3(\text{empty}))) : \text{SORTonce}$$

□

Horizontal composability for constructor implementations takes the form:

$$\frac{SP_1 \rightsquigarrow_{K_1'} SP_1' \quad \cdots \quad SP_n \rightsquigarrow_{K_n'} SP_n'}{op(SP_1, \dots, SP_n) \rightsquigarrow op(\bar{K}_1'(SP_1'), \dots, \bar{K}_n'(SP_n'))}$$

The problem illustrated by Example 2 still exists, but it cannot arise when *op* corresponds to a parameterized program, as in the decomposition steps via multi-argument parameterized programs above.

8. Specifying and Developing Parameterized Programs

The enterprise of formal specification and development is relevant to parameterized programs, with exactly the same motivation as in the case of ordinary non-parameterized programs. An additional advantage this brings is that it enables the overall shape of a development tree (see above) to be given without the need to supply the parameterized programs involved in each of the steps. The provision of a parameterized program that fits into each step can then be regarded as a separate task, perhaps involving further refinement and decomposition. This is one of the ideas underlying the Extended ML methodology for formal development of Standard ML programs from specifications [SaT89, SaT91, San91, Kaz92].

For any two signatures Σ and Σ' , we can regard $\Sigma \rightarrow \Sigma'$ as a new kind of signature. Then $Alg(\Sigma \rightarrow \Sigma')$ is the set of all *parametric* $(\Sigma \rightarrow \Sigma')$ -algebras, that is functions $F : Alg(\Sigma) \rightarrow Alg(\Sigma')$. Just as ordinary programs are modelled as algebras, parameterized programs are modelled as parametric algebras. (We generalize this further to multi-argument and higher-order parametric algebras below.)

Note that both “constructor” and “parametric algebra” are names for the same concept: a function mapping algebras to algebras. We use the former when such a function constitutes an implementation step, and the latter when it is itself the outcome of a development task. This distinction is blurred below, especially once the extension to higher-order is considered. Another difference is that a constructor is assumed to be defined by a parameterized program, while a parametric algebra is an arbitrary set-theoretic function.

To specify a parameterized program, we give its input and output interfaces. The specification $SP \rightarrow SP'$ describes the class of parametric $(Sig(SP) \rightarrow Sig(SP'))$ -algebras $F : Alg(Sig(SP)) \rightarrow Alg(Sig(SP'))$ such that $F(A) \in \llbracket SP' \rrbracket$ for all $A \in \llbracket SP \rrbracket$. Said another way, $Sig(SP \rightarrow SP') = Sig(SP) \rightarrow Sig(SP')$ and $\llbracket SP \rightarrow SP' \rrbracket = \{F \in Alg(Sig(SP) \rightarrow Sig(SP')) \mid F(\llbracket SP \rrbracket) \subseteq \llbracket SP' \rrbracket\}$. The statement that κ is a realization of $SP \rightarrow SP'$ is thus equivalent to the correctness of the constructor implementation $SP' \rightsquigarrow_{\kappa} SP$. The specification $SP \rightarrow SP'$ is *not* a so-called *parameterized specification*; it is a non-parameterized specification of a parameterized program. See [SST92] for a discussion of this distinction.

Example 1 (continued). $INTLIST \rightarrow SORT$ specifies a parameterized program which, given an implementation of $INTLIST$, delivers an implementation of $SORT$. Two (equivalent) realizations of this specification are the functors:

```
functor K(X:INTLIST):SORT = K0(K1(K2(X)))
functor K'(X:INTLIST):SORT = K1(K2(X))
```

Another, different realization is:

```
functor K''(X:INTLIST):SORT = K4(K1(K2(X)),K5(X))
```

That is, $\llbracket K \rrbracket, \llbracket K' \rrbracket, \llbracket K'' \rrbracket \in \llbracket INTLIST \rightarrow SORT \rrbracket$. \square

The definition of refinement of specifications applies without modification to specifications of parameterized programs:

$$SP_1 \rightarrow SP'_1 \rightsquigarrow SP_2 \rightarrow SP'_2 \quad \text{iff} \quad \llbracket SP_2 \rightarrow SP'_2 \rrbracket \subseteq \llbracket SP_1 \rightarrow SP'_1 \rrbracket$$

which again presupposes that $Sig(SP_1 \rightarrow SP'_1) = Sig(SP_2 \rightarrow SP'_2)$, i.e. $Sig(SP_1) = Sig(SP_2)$ and $Sig(SP'_1) = Sig(SP'_2)$. A sufficient condition for this refinement to hold is that $SP_2 \rightsquigarrow SP_1$ and $SP'_1 \rightsquigarrow SP'_2$.

Example 1 (continued). Simple examples of refinements between specifications of parameterized programs may be built on the examples of refinements given in Example 1 in Section 6: $INTLIST \rightarrow SORT$ refines to $INTLIST \rightarrow SORT_{perm}$ which further refines to $INTLIST \rightarrow SORT_{ins}$ which refines to $INTLIST \rightarrow SORT_{done}$. \square

The above presentation uses a particularly simple form of specification of parameterized programs, where the output interface does not depend on the particular realization of the input interface. This is not sufficient when more complex examples are considered. The necessary extra flexibility is gained by replacing the specification $SP \rightarrow SP'$ by the generalized (dependent) product $\Pi X:SP.SP'[X]$. See [SST92] for details of this and other technicalities omitted here.

Example 1 (continued). Specifications like $INTLIST \rightarrow SORT$ do not capture the intention that their realizations, when given an argument X realizing $INTLIST$, should produce a realization of $SORT$ that *extends* X . So a realization of $INTLIST \rightarrow SORT$ might ignore the *po* component of its argument and produce a realization of $SORT$ containing a completely different *po* function, together with a sorting function that is correct with respect to this new *po* function rather than the one supplied in the argument. This problem can be solved by use of the following dependent product specification:

$\Pi X:INTLIST$.
enrich $SORT$ **by**
axioms $\forall x, y: \text{int}. po(x, y) = X.po(x, y)$
 $\forall x: \text{int}. \forall l: \text{int list}. is_in(x, l) = X.is_in(x, l)$
 $\forall l: \text{int list}. head(l) = X.head(l) \wedge tail(l) = X.tail(l)$

In case types are involved, this is the issue of *sharing* in Standard ML and the use of so-called *sharing constraints* as in Standard ML [Pau91] and Extended ML [SaT89, SaT91] is one way of expressing the required dependency. See SPECTRAL [KBS91] for a different approach.

The need to “copy” the components of X (*po*, *is_in*, etc.) may seem ugly. In fact, since these components are provided by the argument X , there is no need to include them explicitly in the result — if they are needed later on somewhere else, they can always be recovered directly from X itself rather than via the result. This would lead to the following dependent product specification:

$\Pi X:INTLIST$.
hide opns is_sorted **in**
opns $is_sorted : \text{int list} \rightarrow \text{bool}$
 $sort : \text{int list} \rightarrow \text{int list}$
axioms $is_sorted(\text{nil}) = \text{true}$
 $\forall x: \text{int}. \forall l: \text{int list}.$
 $is_sorted(x::l) = \text{true} \iff$
 $((\forall y: \text{int}. X.is_in(y, l) = \text{true} \implies X.po(x, y) = \text{true})$
 $\wedge is_sorted(l) = \text{true})$
 $\forall l: \text{int list}. is_sorted(sort(l)) = \text{true}$
 $\forall l: \text{int list}. \forall x: \text{int}. X.is_in(x, l) = X.is_in(x, sort(l))$

\square

Constructor implementations may be similarly generalized to deal with specifications of parameterized programs. However, the parameterized programs used to

define the constructors involved in such implementations are then higher-order, i.e. they take parameterized programs as arguments and return parameterized programs as results. Higher-order functors are not available in Standard ML as defined in [MTH90], but their semantics and implementation is a topic of current active research [Tof92, MaT94, Bis95].

8.1. Higher-Order Parameterization

The definitions involved in dealing with parameterized programs and their specifications extend to the higher-order case in a natural way [SST92]. The set of generalized signatures is defined to be the least set containing ordinary signatures and such that if $\Sigma_1, \dots, \Sigma_n$ ($n \geq 0$) and Σ are generalized signatures, then $\langle \Sigma_1, \dots, \Sigma_n \rangle \rightarrow \Sigma$ is a generalized signature as well; if $n = 1$ we omit the brackets. $Alg(\langle \Sigma_1, \dots, \Sigma_n \rangle \rightarrow \Sigma)$ is the set of all functions $F : Alg(\Sigma_1) \times \dots \times Alg(\Sigma_n) \rightarrow Alg(\Sigma)$. (Alternatively, multi-argument parameterized programs could be presented in their “curried” form denoting parametric algebras in $Alg(\Sigma_1 \rightarrow (\Sigma_2 \rightarrow \dots (\Sigma_n \rightarrow \Sigma) \dots))$.) As before, to specify a higher-order parameterized program, we give its input interfaces and output interface, which may now themselves be specifications of (higher-order) parameterized programs. The definitions are exactly the same as those given above. In the following, for simplicity, we omit the problems of dependency of the result specification on the arguments; as before, a solution is to use generalized product specifications.

The concepts of constructor implementation and decomposition step carry over without modification to the case of higher-order parameterized programs and their specifications.

Example 1 (continued). The example of constructor implementation with decomposition at the end of the last section may be rephrased using these ideas. Instead of implementing *SORTonce* in terms of realizations of *SORTperm* and *NOSTUTTER*, we can build a realization of the specification $INTLIST \rightarrow SORTonce$ in terms of realizations of the specifications $INTLIST \rightarrow SORTperm$ and $INTLIST \rightarrow NOSTUTTER$:

$$INTLIST \rightarrow SORTonce \rightsquigarrow_{K6} \langle INTLIST \rightarrow SORTperm, INTLIST \rightarrow NOSTUTTER \rangle$$

where $K6$ is the following higher-order functor (written using an *ad hoc* but hopefully self-explanatory notation):

```

functor K6( F1:INTLIST->SORTperm, F2:INTLIST->NOSTUTTER )
: INTLIST->SORTonce =
  ( functor (X:INTLIST):SORTonce =
    struct
      structure S = F1(X)
      open S
      structure N = F2(X)
      fun sort l = N.rem_stutter(S.sort(l))
    end )

```

□

It is possible to restrict attention to parameterized programs of a particularly simple form, since any constructor implementation $SP \rightsquigarrow_K \langle SP_1, \dots, SP_n \rangle$

may be replaced by the decomposition $SP \xrightarrow{\sim}_{app} \langle SP_\kappa, SP_1, \dots, SP_n \rangle$, where $SP_\kappa = \langle SP_1, \dots, SP_n \rangle \rightarrow SP$ and app is the higher-order parameterized program such that $app(F, A_1, \dots, A_n) = F(A_1, \dots, A_n)$ and where the parameterized program κ is then provided as the realization of SP_κ . A decomposition like $SP \xrightarrow{\sim}_{app} \langle SP_\kappa, SP_1, \dots, SP_n \rangle$ embodies the decision to implement SP in terms of realizations of SP_1, \dots, SP_n , leaving the decision of how these are used to produce a realization of SP as a *separate* development task, specified by SP_κ . This brings a bottom-up flavour into our principally top-down view of the development process.

Example 1 (continued). Here is the above example once again:

$$\begin{aligned} &INTLIST \rightarrow SORT_{once} \xrightarrow{\sim}_{APP} \\ &\quad \langle \langle INTLIST \rightarrow SORT_{perm}, INTLIST \rightarrow NOSTUTTER \rangle \rightarrow \\ &\quad \quad \quad (INTLIST \rightarrow SORT_{once}), \\ &\quad INTLIST \rightarrow SORT_{perm}, \\ &\quad INTLIST \rightarrow NOSTUTTER \rangle \end{aligned}$$

where APP is the higher-order functor applying its first argument to its second and third arguments:

```
functor APP
  ( F: (INTLIST->SORTperm, INTLIST->NOSTUTTER) ->
    (INTLIST->SORTonce),
    F1: INTLIST->SORTperm, F2: INTLIST->NOSTUTTER
  ): INTLIST->SORTonce =
  F(F1, F2)
```

This embodies a decision that the implementation of $INTLIST \rightarrow SORT_{once}$ may use the implementations of the specifications $INTLIST \rightarrow SORT_{perm}$ and $INTLIST \rightarrow NOSTUTTER$, to be provided separately. One way of realizing the specification

$$\langle \langle INTLIST \rightarrow SORT_{perm}, INTLIST \rightarrow NOSTUTTER \rangle \rightarrow (INTLIST \rightarrow SORT_{once})$$

is the functor K6 above; another possibility is to use an entirely different solution, ignoring either or both of the realizations of $INTLIST \rightarrow SORT_{perm}$ and $INTLIST \rightarrow NOSTUTTER$. For example:

```
functor K7( F1: INTLIST->SORTperm, F2: INTLIST->NOSTUTTER )
: INTLIST->SORTonce =
  ( functor (X: INTLIST): SORTonce =
    struct
      open X
      fun insert(x, nil) = [x]
      | insert(x, y::l) =
        if x=y then y::l
        else if po(x,y) then x::y::l
        else y::insert(x,l)
      fun sort(nil) = nil
      | sort(x::l) = insert(x, sort(l))
    end )
```

□

9. Behavioural Implementations

A specification should be a precise and complete statement of required properties. We should try to avoid including extra requirements, even if they happen to be satisfied by a possible future realization. Such over-specification unnecessarily limits the options left open to the implementer. Ideally, the target is to describe exactly the admissible program behaviours. This suggests that specifications of programming tasks should not distinguish between programs (modelled as algebras) exhibiting the same *behaviour*.

The intuitive idea of behaviour of an algebra has been formalised in a number of ways (see e.g. [Rei81, GoM82, SaW83, Sch87, SaT87, NiO88]).⁵ In most approaches one distinguishes a certain set *OBS* of sorts as *observable*. Intuitively, these are the sorts of data directly visible to the user (integers, booleans, characters, etc.) in contrast to sorts of “internal” data structures, which are observable only via the functions provided by the program. The behaviour of an algebra is characterised by the set of *observable computations* taking arguments of sorts in *OBS* and producing a result of a sort in *OBS*. In the standard algebraic framework, such computations are modelled as terms of sorts in *OBS* with variables (representing the inputs) of sorts in *OBS* only. Two Σ -algebras *A* and *B* are *behaviourally equivalent* (w.r.t. *OBS*), written $A \equiv B$, if they exhibit the same behaviour, that is, if all observable computations yield the same results in *A* and in *B*. The motivation is related to that of so-called *testing* equivalences studied in the context of concurrent systems [DNH84]. The role of behavioural equivalence in the context of parametric algebras is a topic of current research and we do not treat this here. Therefore this section deals only with ordinary algebras and development of non-parameterized programs.

Example 4. A hackneyed example that illustrates the idea of behavioural equivalence is that of stacks of integers:

STACK	=	sorts	<i>stack</i>
		opns	<i>empty</i> : <i>stack</i> → bool <i>push</i> : int × <i>stack</i> → <i>stack</i> <i>pop</i> : <i>stack</i> → <i>stack</i> <i>top</i> : <i>stack</i> → int <i>is_empty</i> : <i>stack</i> → bool
		axioms	<i>is_empty</i> (<i>empty</i>) = true $\forall s:\textit{stack}. \forall n:\textit{int}. \textit{is_empty}(\textit{push}(n, s)) = \textit{false}$ $\forall s:\textit{stack}. \forall n:\textit{int}. \textit{top}(\textit{push}(n, s)) = n$ $\forall s:\textit{stack}. \forall n:\textit{int}. \textit{pop}(\textit{push}(n, s)) = s$

Suppose that the sorts *int* and *bool* (included implicitly in all the specifications we consider) are taken as observable while the sort *stack* is not. The observable computations are all the terms of the form *is_empty*(*s*) and *top*(*s*) where *s* is a term of sort *stack* with variables of sort *int* only.

Two typical algebras which provide intuitively acceptable realizations of this specification can be coded as SML structures as follows:

⁵ The following paragraph makes sense only in an institution in which signatures have sorts. This is not much of a restriction in practice. In any case, what follows thereafter (apart from the examples) applies to any institution and any equivalence relation on its algebras [SaT88b].


```

structure S1:STACK =
  struct
    type stack = int list
    val empty = nil
    fun push(n,s) = n::s
    fun pop(nil) = nil
      | pop(_::s) = s
    fun top(nil) = 0
      | top(n::_) = n
    fun is_empty(nil) = true
      | is_empty(_::_) = false
  end

structure S2:STACK =
  struct
    type stack = (int -> int) * int
    val empty = ((fn k => 0), 0)
    fun push(n,(f,i)) =
      ((fn k => if k = i then n else f k), i+1)
    fun pop(f,i) = if i = 0 then (f,0) else (f,i-1)
    fun top(f,i) = if i = 0 then 0 else f(i-1)
    fun is_empty(f,i) = (i=0)
  end

```

S1 gives the obvious realization of stacks as list of integers and S2 codes the realization of stacks as (infinite) arrays with pointers to the top of the represented stack (arrays are coded here as functions from integer indices to values).

Now, these two realizations of stacks are behaviourally equivalent since for each observable computation, like $top(pop(push(n, push(4, push(6, empty))))$, they both deliver the same result (in this case 4). However, these algebras do not act the same way when non-observable computations are considered: for example, the computations $empty$ and $pop(push(6, empty))$ yield the same result in S1 but they yield different results in S2. \square

Our earlier discussion would lead us to expect the class of models of a specification to be closed under behavioural equivalence. It is perhaps surprising that this is not easy to achieve directly: the class of models of a set of axioms typically does not have this property. Equational logic may be modified so as to force this to happen (cf. [NiO88]) and a similar idea for other logical systems is discussed in [BHW95], but it is not clear how this approach can be extended to deal adequately with structured specifications. An alternative is to simply close the class of models of a specification under behavioural equivalence [SaW83, SaT87]. Any specification SP determines the class $\llbracket SP \rrbracket \subseteq Alg(Sig(SP))$ of models that “literally” satisfy the stated requirements, as discussed in Section 3; the ultimate semantics of SP is taken to be the closure of this under behavioural equivalence:

$$\llbracket\llbracket SP \rrbracket\rrbracket = \{A \mid A \equiv B \text{ for some } B \in \llbracket SP \rrbracket\}$$

In particular, $\llbracket\text{sorts } S \text{ opns } \Omega \text{ axioms } \Phi\rrbracket$ contains exactly the $\langle S, \Omega \rangle$ -algebras that satisfy the axioms Φ , while $\llbracket\llbracket\text{sorts } S \text{ opns } \Omega \text{ axioms } \Phi\rrbracket\rrbracket$ contains also the algebras that do not satisfy Φ themselves but are behaviourally equivalent to algebras that do. The notation $\llbracket\llbracket\cdot\rrbracket\rrbracket$ applies to specifications only; it does not apply to programs or parameterized program.

Example 4 (continued). Both the list representation $S1$ and the array-with-pointer representation $S2$ of stacks are in $\llbracket \text{STACK} \rrbracket$ — and this is what we meant when we declared

```
structure S1:STACK = ...
structure S2:STACK = ...
```

That is, $S1, S2 \in \llbracket \text{STACK} \rrbracket$. This holds even though the latter realization $S2$ does not literally satisfy the axiom $\forall s:\text{stack}. \forall n:\text{int}. \text{pop}(\text{push}(n, s)) = s$ and so $S2 \notin \llbracket \text{STACK} \rrbracket$. \square

This approach typically gives extra expressive power: considering the institution of first-order logic with equality, there are classes of algebras that may be finitely characterized in this way, and cannot be finitely axiomatized directly [Sch92]. (Of course, this property depends on the logic considered: for example, second-order logic allows one to specify behavioural closures directly [HoS96].) Also, *model-oriented specifications* [Jon86] can be handled: if $\llbracket SP \rrbracket$ contains just a single algebra, then $\llbracket SP \rrbracket$ admits any realization of the exhibited behaviour. In general, $\llbracket SP \rrbracket$ contains all *reifications* of the algebras in $\llbracket SP \rrbracket$ (cf. [Hoa72]).

The basic intuition for the use of behavioural equivalence in the development process is that it is not necessary to implement a specification SP according to its literal interpretation $\llbracket SP \rrbracket$; it is sufficient to implement it up to behavioural equivalence, as captured by its “ultimate” semantics $\llbracket SP \rrbracket$. The definitions of refinement and constructor implementation are now as follows:

$$\begin{aligned} SP \rightsquigarrow SP' & \text{ iff } \llbracket SP' \rrbracket \subseteq \llbracket SP \rrbracket, \\ SP \rightsquigarrow_{\kappa} SP' & \text{ iff } \llbracket \kappa \rrbracket(\llbracket SP' \rrbracket) \subseteq \llbracket SP \rrbracket. \end{aligned}$$

Using these definitions, it is possible to develop programs from specifications by means of successive implementation steps exactly as described earlier.

9.1. Stability

The development process may take advantage of the behavioural interpretation of specifications in a more delicate way than suggested above. The crucial novelty, due to [Sch87], is that when *using* a realization of SP , it is convenient (and possible) to pretend that it satisfies the literal interpretation of SP .

Example 4 (continued). Consider the following trivial specification

```
TRIV =  ops    id : int × int × int → int
        axioms  ∀x, n, z:int. n >= 0 ⇒ id(x, n, z) = x
```

and its perhaps surprising realization in terms of stacks of integers:⁶

```
functor TR(S:STACK):TRIV =
  struct
    fun multipush(n,z,s:S.stack) =
      if n <= 0 then s
      else S.push(z,multipush(n-1,z+1,s))
```

⁶ This is of course an extremely contrived example, but it is easy to come up with realistic programs using stacks where properties like this are to be proved.

```

fun multipop(n,s:S.stack) =
  if n <= 0 then s else multipop(n-1,S.pop(s))
fun id(x,n,z) =
  S.top(multipop(n,multipush(n,z,
                           S.push(x,S.empty))))
end

```

Now, given any realization S of $STACK$, to verify that $TR(S) \in \llbracket TRIV \rrbracket$, it is convenient to assume that the axiom $\forall s:stack. \forall n:int. pop(push(n,s)) = s$ holds in S literally in spite of the fact that this equation is not valid in $\llbracket STACK \rrbracket$. Under this assumption, a simple proof by induction (on the second argument of id) goes through. The reasoning for the induction step goes as follows:

```

id(x,n+1,z)
= S.top(multipop(n+1,multipush(n+1,z,S.push(x,S.empty))))
= S.top(multipop(n,S.pop(S.push(z,multipush(n,z+1,
                                             S.push(x,S.empty)))))
= S.top(multipop(n,multipush(n,z+1,S.push(x,S.empty))))
= id(x,n,z+1)
= x

```

where the final step follows by the induction hypothesis. \square

These considerations lead to the following definition of *behavioural implementation* [SaT88b]:

$$SP \xrightarrow{\kappa} SP' \text{ iff } \llbracket \kappa \rrbracket(\llbracket SP' \rrbracket) \subseteq \llbracket SP \rrbracket$$

The alert reader will have noticed that there is a problem here: we want to have our cake and eat it. On one hand, we want to allow specifications to be implemented up to behavioural equivalence; on the other hand, we would like to use any realization as if it satisfied its specification literally. Behavioural implementations do not compose, and the following crucial property is lost:

$$\frac{SP \xrightarrow{\kappa} SP' \quad A' \in \llbracket SP' \rrbracket}{\llbracket \kappa \rrbracket(A') \in \llbracket SP \rrbracket}$$

The behavioural implementation $SP \xrightarrow{\kappa} SP'$ ensures only that algebras in $\llbracket SP' \rrbracket$ give rise to correct realizations of SP ; this says nothing about the models in $\llbracket SP' \rrbracket$ that are not in $\llbracket SP' \rrbracket$.

Example 4 (continued). As stated above, $TRIV \not\sim_{TR} STACK$. Formally, neither the property this embodies (that $TR(S) \in \llbracket TRIV \rrbracket$ for $S \in \llbracket STACK \rrbracket$) nor its suggested proof tell us anything about the application of TR to the algebra $S2 \notin \llbracket STACK \rrbracket$, even though $S2 \in \llbracket STACK \rrbracket$ and we have earlier identified $S2$ as an acceptable realization of $STACK$. It may be shown that $TR(S2) \in \llbracket TRIV \rrbracket$ but the most obvious proof involves the non-elementary fact that for any natural number n ,

$$multipop(n,s) = \underbrace{S.pop(\dots(S.pop(s))\dots)}_{n \text{ times}}$$

and similarly for $multipush$ (and then relies on the property that since $S2 \in \llbracket STACK \rrbracket$, all observable computations in $S2$ yield the same results on s and on $pop(push(z,s))$, for any stack s and integer z).

Consider, however, another trivial realization of $TRIV$ in terms of $STACK$:

```

functor TR' (S:STACK):TRIV =
  struct
    fun id(x,n,z) = let val se = S.pop(S.push(z,S.empty))
                     in if se = S.empty then x else z
                     end
  end

```

(TR' cannot be coded in SML since the type `S.stack` is not ensured to admit equality.)

We can prove now that $TRIV \stackrel{\text{TR}'}{\sim} STACK$, hence $TR'(S1) \in \llbracket TRIV \rrbracket$, but of course $TR'(S2) \notin \llbracket TRIV \rrbracket$. \square

It might seem that all is lost. But there is a way out, originally suggested in [Sch87]. The above crucial property is recovered if we assume that the constructors used are *stable*, that is, that any constructor $\llbracket \kappa \rrbracket : Alg(Sig(SP')) \rightarrow Alg(Sig(SP))$ preserves behavioural equivalence:

Stability assumption: if $A \equiv B$ then $\llbracket \kappa \rrbracket(A) \equiv \llbracket \kappa \rrbracket(B)$

(the exact definition of stability of constructors in a formal development framework based on a full-blown programming language is somewhat more complex — see [Sch87, SaT89]).

Under this assumption, the correctness of the individual implementation steps ensures the correctness of the result:

$$\frac{SP_0 \stackrel{\text{TR}'}{\sim}_{\kappa_1} SP_1 \stackrel{\text{TR}'}{\sim}_{\kappa_2} \cdots \stackrel{\text{TR}'}{\sim}_{\kappa_n} SP_n = EMPTY}{\llbracket \kappa_1(\kappa_2(\dots \kappa_n(empty) \dots)) \rrbracket \in \llbracket SP_0 \rrbracket}$$

Example 4 (continued). Clearly, the functor TR' as defined above is not stable.

On the other hand, the functor TR is stable (this can be proved along the lines of the argument given above to justify that $TR(S2) \in \llbracket TRIV \rrbracket$). This shows that $TR(S) \in \llbracket TRIV \rrbracket$ for all $S \in \llbracket STACK \rrbracket$, not only for $S \in \llbracket STACK \rrbracket$. \square

We could repeat here the tree-like development picture of Section 7 — developments involving decomposition steps based on behavioural implementations with multi-argument (stable) constructors yield correct programs as well. We also recover vertical composability, under the assumption that parameterized programs compose as discussed in Section 7:

$$\frac{SP \stackrel{\text{TR}'}{\sim}_{\kappa} SP' \quad SP' \stackrel{\text{TR}'}{\sim}_{\kappa} SP''}{SP \stackrel{\text{TR}'}{\sim}_{\kappa; \kappa} SP''}$$

The correctness of a behavioural implementation $SP \stackrel{\text{TR}'}{\sim}_{\kappa} SP'$ is easier to verify than the correctness of the corresponding constructor implementation between the same specifications closed under behavioural equivalence: the condition $\llbracket \kappa \rrbracket(\llbracket SP' \rrbracket) \subseteq \llbracket SP \rrbracket$ is weaker than $\llbracket \kappa \rrbracket(\llbracket SP' \rrbracket) \subseteq \llbracket SP \rrbracket$ (but semantically these two conditions become equivalent if $\llbracket \kappa \rrbracket$ is stable). Also, the correctness of $SP \stackrel{\text{TR}'}{\sim}_{\kappa} SP'$ is in general easier to verify than the correctness of the original constructor implementation $SP \stackrel{\text{TR}'}{\sim}_{\kappa} SP'$ (that is, $\llbracket \kappa \rrbracket(\llbracket SP' \rrbracket) \subseteq \llbracket SP \rrbracket$ is weaker than $\kappa(\llbracket SP' \rrbracket) \subseteq \llbracket SP \rrbracket$). For instance, viewing the structure S2 of Example 4 as a functor with an empty parameter, we have $STACK \stackrel{\text{TR}'}{\sim}_{S2} EMPTY$, while $STACK \stackrel{\text{TR}'}{\sim}_{S2} EMPTY$. As we have argued, this extra flexibility reflects our intuitive understanding of what it means for an algebra to realize a specification.

We are still left with the need to establish the stability of constructors, and so one may wonder if it is worthwhile taking advantage of this property. However, recall that constructors are determined by parameterized programs, and these must be expressed in some particular programming language. Thus stability can be checked in advance, for the programming language as a whole (this is simplified somewhat by the fact that the composition of stable constructors is stable) and this frees the programmer from the need to prove it during the program development process. Other views of stability are possible, cf. [NOS95].

There is a close connection between the requirement of stability and the security of encapsulation mechanisms in programming languages supporting abstract data types. A programming language ensures stability if the only way to access an encapsulated data type is via the operations explicitly provided in its output interface. This suggests that stability of constructors is an appropriate thing to expect; following [Sch87] we view the stability requirement as a methodologically justified design criterion for the modularization facilities of programming languages.

Example 5. The limited notation of Standard ML functors we have used in Examples 1 and 4 throughout this paper ensures stability of constructors.⁷ Intuitively, this is because the parameter signature of any SML functor provides sufficient insulation between the functor body and the actual parameters. Within the functor body we can access the parameter only using the “tools” given in the parameter signature. Unfortunately, signatures given for functor results and for structures declared in SML are much more “transparent”: they do not provide sufficient insulation between the declaration of a structure and its use. For example, the following code in SML shows a non-stable extension of the structure S2:

```
structure S2:STACK =
  struct
    type stack = (int -> int) * int
    val empty = ...
    ...
    fun pop(f,i) = if i = 0 then (f,0) else (f,i-1)
    ...
  end
fun id(x,n,z) = let val (f,i) = S2.pop(S2.push(z,S2.empty))
                in if f i = 0 then x else z
                end
```

Changing the “internal” implementation details of S2 for example as follows:

```
structure S2:STACK =
  struct
    type stack = (int -> int) * int
    val empty = ...
    ...
    fun pop(f,i) =
      if i = 0 then (f,0)
```

⁷ Of course, considerable work would be required to turn this claim into a formal theorem with a precise proof.

```

else (fn k => if k = i-1 then 0 else f k, i-1)
...
end

```

changes completely the behaviour of the `id` function as defined in the extension, even though the new realization `S2` of stacks is behaviourally equivalent to the previous one.

This can be viewed as a deficiency in the design of the Standard ML modularization facilities. This infelicity is not present in the Extended ML formalism [SaT91, KST94, KST97] where access to a structure or functor result is limited to the use of the tools given in its signature. \square

10. Conclusion

We have outlined the main ideas of a framework to support the formal development of correct programs from specifications of their required behaviour. Our purpose has not been to introduce new technicalities, but rather to explain in a careful way the general ideas underlying the algebraic approach and the specific motivation behind the concepts involved in the formalization of the development process. This forced us to clarify some of the finer points of the approach, like the distinction between syntax and semantics in constructor implementations and an abstract formulation of stability in this context.

The main challenge now is to put these ideas into practice in the formal development of non-trivial programs in real programming languages. We are moving in this direction with our work on the Extended ML framework for the formal development of modular Standard ML programs [SaT91, KST94, KST97], although more effort is required. Subjecting foundational work to the test of practice is sure to bring fascinating new problems and issues to light.

Acknowledgements

Most of the above ideas have been presented elsewhere, and have been discussed with and influenced by many of our colleagues. Thanks especially to Michel Bidoit, Rod Burstall, Jordi Farrés, Joseph Goguen, Fernando Orejas, Oliver Schoett and Martin Wirsing. Thanks also to Bernd Krieg-Brückner, Luis Dominguez and the referees for comments which helped to improve the presentation. This research was supported by the EC-funded COMPASS Basic Research working group and MeDiCiS and EuroFoCS Scientific Cooperation Networks (DS, AT), by EPSRC grants GR/H73103 and GR/J07303 and an EPSRC Advanced Fellowship (DS), and by KBN grant 2 P301 007 04 (AT).

References

- [AsR95] Astesiano, E. and Reggio, G.: Formally-driven friendly specifications of concurrent systems: a two-rail approach. *Proc. ICSE-17 Workshop on Formal Methods Application in Software Engineering Practice*, Seattle, 158–165 (1995).
- [Bar74] Barwise, J.: Axioms for abstract model theory. *Ann. Math. Logic* 7:221–265 (1974).
- [BaF85] Barwise, J. and Feferman, E.: (eds.) *Model-Theoretic Logics*. Springer (1985).
- [BaW82] Bauer, F. and Wössner, H.: *Algorithmic Language and Program Development*. Springer (1982).

- [BeV87] Beierle, C. and Voß, A.: Viewing implementations as an institution. *Proc. 1987 Conference on Category Theory and Computer Science*, Edinburgh. Springer LNCS 283, 196–218 (1987).
- [BGM89] Bidoit, M., Gaudel, M.-C. and Mauboussin, A.: How to make algebraic specifications more understandable? An experiment with the PLUSS specification language. *Science of Computer Programming* 12:1–38 (1989).
- [BHW95] Bidoit, M., Hennicker, R. and Wirsing, M.: Behavioural and abstractor specifications. *Science of Computer Programming* 25:149–186 (1995).
- [BKL91] Bidoit, M., Kreowski, H.-J., Lescanne, P., Orejas, F. and Sannella, D.: (eds.) *Algebraic System Specification and Development: A Survey and Annotated Bibliography*. Springer LNCS 501 (1991).
- [Bis95] Biswas, S.: Higher-order functors with transparent signatures. *Proc. 22nd ACM Symp. on Principles of Programming Languages*, San Francisco, 154–163 (1995).
- [BuG77] Burstall, R. and Goguen, J.: Putting theories together to make specifications. *Proc. 5th Intl. Joint Conf. on Artificial Intelligence*, Cambridge, Massachusetts, 1045–1058 (1977).
- [BuG80] Burstall, R. and Goguen, J.: The semantics of Clear, a specification language. *Proc. 1979 Copenhagen Winter School on Abstract Software Specification*. Springer LNCS 86, 292–332 (1980).
- [Cen94] Cengarle, M.V.: Formal specifications with higher-order parameterization. Ph.D. thesis, LMU München (1994).
- [Dav90] Davis, A.: *Software Requirements: Analysis and Specification*. Prentice Hall (1990).
- [DNH84] De Nicola, R. and Hennessy, M.: Testing equivalences for processes. *Theoretical Computer Science* 34:83–133 (1984).
- [DGS93] Diaconescu, R., Goguen, J. and Stefanescu, P.: Logical support for modularization. In: *Logical Environments* (G. Huet and G. Plotkin, eds.). Cambridge Univ. Press, 83–130 (1993).
- [EBO93] Ehrig, H., Baldamus, M. and Orejas, F.: New concepts of amalgamation and extension of a general theory of specifications. *Selected Papers from the 8th Workshop on Specification of Abstract Data Types*, Dourdan. Springer LNCS 655, 199–221 (1993).
- [EKM82] Ehrig, H., Kreowski, H.-J., Mahr, B. and Padawitz, P.: Algebraic implementation of abstract data types. *Theoretical Computer Science* 20:209–263 (1982).
- [EhM85] Ehrig, H. and Mahr, B.: *Fundamentals of Algebraic Specification I: Equations and Initial Semantics*. Springer (1985).
- [Far92] Farrés-Casals, J.: Verification in ASL and Related Specification Languages. Ph.D. thesis, report CST-92-92, Dept. of Computer Science, Univ. of Edinburgh (1992).
- [FiS88] Fiadeiro, J. and Sernadas, A.: Structuring theories on consequence. *Selected Papers from the 5th Workshop on Specification of Abstract Data Types*, Gullane. Springer LNCS 332, 44–72 (1988).
- [FiJ90] Fitzgerald, J. and Jones, C.: Modularizing the formal description of a database system. *Proc. VDM'90 Conference*, Kiel. Springer LNCS 428, 198–210 (1990).
- [Gog84] Goguen, J.: Parameterized programming. *IEEE Trans. on Software Engineering* SE-10(5):528–543 (1984).
- [GoB80] Goguen, J. and Burstall, R.: CAT, a system for the structured elaboration of correct programs from structured specifications. Technical report CSL-118, SRI International (1980).
- [GoB84] Goguen, J. and Burstall, R.: Introducing institutions. *Proc. Logics of Programming Workshop*, Carnegie-Mellon. Springer LNCS 164, 221–256 (1984).
- [GoB92] Goguen, J. and Burstall, R.: Institutions: abstract model theory for specification and programming. *Journal of the Assoc. for Computing Machinery* 39:95–146 (1992).
- [GoL95] Goguen, J. and Luqi.: Formal methods and social context in software development. *Proc. 6th Joint Conf. on Theory and Practice of Software Development*, TAPSOFT'95, Aarhus. Springer LNCS 915, 62–81 (1995).
- [GoM82] Goguen, J. and Meseguer, J.: Universal realization, persistent interconnection and implementation of abstract modules. *Proc. Intl. Colloq. on Automata, Languages and Programming*, Aarhus. Springer LNCS 140, 265–281 (1982).
- [GoM85] Goguen, J. and Meseguer, J.: Completeness of many-sorted equational logic. *Houston Journal of Mathematics* 11(3):307–334 (1985).
- [GWM92] Goguen, J., Winkler, T., Meseguer, J., Futatsugi, K. and Jouannaud, J.-P.: Introducing OBJ3. Technical Report SRI-CSL-92-03, SRI International (1992).
- [GuH80] Guttag, J. and Horning, J.: Formal specification as a design tool. *Proc. 7th ACM Symp. on Principles of Programming Languages*, Las Vegas, 251–261 (1980).

- [GuH93] Guttag, J. and Horning, J.: *Larch: Languages and Tools for Formal Specification*. Springer (1993).
- [GHW82] Guttag, J., Horning, J. and Wing, J.: Some notes on putting formal specifications to productive use. *Science of Computer Programming* 2:53–68 (1982).
- [HST94] Harper, R., Sannella, D. and Tarlecki, A.: Structured theory presentations and logic representations. *Annals of Pure and Applied Logic* 67:113–160 (1994).
- [HaJ89] Hayes, I.J. and Jones, C.B.: Specifications are not (necessarily) executable. *Software Engineering Journal* 4(6):320–338 (1989).
- [Hoa72] Hoare, C.A.R.: Proofs of correctness of data representations. *Acta Informatica* 1:271–281 (1972).
- [HKB93] Hoffmann, B. and Krieg-Brückner, B.: (eds.) *PROgram Development by SPECification and TRANSformation: Methodology – Language Family – System*. Springer LNCS 680 (1993).
- [HoS96] Hofmann, M. and Sannella, D.: On behavioural abstraction and behavioural satisfaction in higher-order logic. *Theoretical Computer Science* 167:3–45 (1996).
- [Huß85] Hußmann, H.: Rapid prototyping for algebraic specifications: RAP system user's manual. Report MIP-8504, Universität Passau (1985).
- [Jon86] Jones, C.: *Systematic Software Development using VDM*. Prentice Hall (1986).
- [KST94] Kahrs, S., Sannella, D. and Tarlecki, A.: The definition of Extended ML. Report ECS-LFCS-94-300, Dept. of Computer Science, Univ. of Edinburgh (1994).
- [KST97] Kahrs, S., Sannella, D. and Tarlecki, A.: The definition of Extended ML: a gentle introduction. *Theoretical Computer Science* 173, to appear (1997).
- [Kaz92] Kazmierczak, E.: Modularizing the specification of a small database system in Extended ML. *Formal Aspects of Computing* 4:100–142 (1992).
- [KBS91] Krieg-Brückner, B. and Sannella, D.: Structuring specifications in-the-large and in-the-small: higher-order functions, dependent types and inheritance in SPECTRAL. *Proc. 4th Joint Conf. on Theory and Practice of Software Development, TAPSOFT'91*, Brighton. Springer LNCS 494, 313–336 (1991).
- [Li94] Li, W.: A logical framework for evolution of specifications. *Proc. 5th European Symp. on Programming*, Edinburgh. Springer LNCS 788, 394–408 (1994).
- [MaT94] MacQueen, D. and Tofte, M.: A semantics for higher-order functors. *Proc. 5th European Symp. on Programming*, Edinburgh. Springer LNCS 788, 409–423 (1994).
- [MTH90] Milner, R., Tofte, M. and Harper, R.: *The Definition of Standard ML*. MIT Press (1990).
- [MoA91] Morris, J. and Ahmed, S.: Designing and refining specifications with modules. *Proc. 3rd Refinement Workshop*, Hursley Park, 1990. Springer Workshops in Computing, 73–95 (1991).
- [NOS95] Navarro, M., Orejas, F. and Sanchez, A.: On the correctness of modular systems. *Theoretical Computer Science* 140:139–177 (1995).
- [Nel91] Nelson, G.: (ed.) *System Programming with Modula-3*. Prentice Hall (1991).
- [NiO88] Nivela, P. and Orejas, F.: Initial behaviour semantics for algebraic specifications. *Selected Papers from the 5th Workshop on Specification of Abstract Data Types*, Gullane. Springer LNCS 332, 184–207 (1988).
- [Ore83] Orejas, F.: Characterizing composability of abstract implementations. *Proc. 1983 Intl. Conf. on Foundations of Computation Theory*, Borgholm. Springer LNCS 158, 335–346 (1983).
- [Pau91] Paulson, L.: *ML for the Working Programmer*. Cambridge Univ. Press (1991).
- [Rei81] Reichel, H.: Behavioural equivalence — a unifying concept for initial and final specification methods. *Proc. 3rd Hungarian Computer Science Conference*, 27–39 (1981).
- [RAH94] Robertson, D., Agustí, J., Hesketh, J. and Levy, J.: Expressing program requirements using refinement lattices. *Fundamenta Informaticae* 21:163–183 (1994).
- [San91] Sannella, D.: Formal program development in Extended ML for the working programmer. *Proc. 3rd BCS/FACS Workshop on Refinement*, Hursley Park. Springer Workshops in Computing, 99–130 (1991).
- [SaB83] Sannella, D. and Burstall, R.: Structured theories in LCF. *Proc. Colloq. on Trees in Algebra and Programming*, L'Aquila. Springer LNCS 159, 377–391 (1983).
- [SST92] Sannella, D., Sokolowski, S. and Tarlecki, A.: Toward formal development of programs from algebraic specifications: parameterization revisited. *Acta Informatica* 29:689–736 (1992).
- [SaT86] Sannella, D. and Tarlecki, A.: Extended ML: an institution-independent framework for formal program development. *Proc. Intl. Workshop on Category Theory and Computer Programming*, Guildford. Springer LNCS 240, 364–389 (1986).

- [SaT87] Sannella, D. and Tarlecki, A.: On observational equivalence and algebraic specification. *J. of Computer and System Sciences* 34:150–178 (1987).
- [SaT88a] Sannella, D. and Tarlecki, A.: Specifications in an arbitrary institution. *Information and Computation* 76:165–210 (1988).
- [SaT88b] Sannella, D. and Tarlecki, A.: Toward formal development of programs from algebraic specifications: implementations revisited. *Acta Informatica* 25:233–281 (1988).
- [SaT89] Sannella, D. and Tarlecki, A.: Toward formal development of ML programs: foundations and methodology. *Proc. 3rd Joint Conf. on Theory and Practice of Software Development*, Barcelona. Springer LNCS 352, 375–389 (1989).
- [SaT91] Sannella, D. and Tarlecki, A.: Extended ML: past, present and future. *Proc. 7th Intl. Workshop on Specification of Abstract Data Types*, Wusterhausen. Springer LNCS 534, 297–322 (1991).
- [SaT9?] Sannella, D. and Tarlecki, A.: *Foundations of Algebraic Specifications and Formal Program Development*. Cambridge Univ. Press, to appear (199?).
- [SaW82] Sannella, D. and Wirsing, M.: Implementation of parameterized specifications. *Proc. Intl. Colloq. on Automata, Languages and Programming*, Aarhus. Springer LNCS 140, 473–488 (1982).
- [SaW83] Sannella, D. and Wirsing, M.: A kernel language for algebraic specification and implementation. *Proc. 1983 Intl. Conf. on Foundations of Computation Theory*, Borgholm. Springer LNCS 158, 413–427 (1983).
- [Sch87] Schoett, O.: Data Abstraction and the Correctness of Modular Programming. Ph.D. thesis, report CST-42-87, Dept. of Computer Science, Univ. of Edinburgh (1987).
- [Sch92] Schoett, O.: Two impossibility theorems on behavioural specification of abstract data types. *Acta Informatica* 29:595–621 (1992).
- [Som92] Sommerville, I.: *Software Engineering* (4th edition). Addison-Wesley (1992).
- [Tof92] Tofte, M.: Principle signatures for higher-order program modules. *Proc. 19th ACM Symp. on Principles of Programming Languages*, Albuquerque, 189–199 (1992).
- [Wir86] Wirsing, M.: Structured algebraic specifications: a kernel language. *Theoretical Computer Science* 42:123–249 (1986).
- [Wir93] Wirsing, M.: Structured specifications: syntax, semantics and proof calculus. *Logic and Algebra of Specification* (F. Bauer, W. Brauer and H. Schwichtenberg, eds.). Springer, 411–442 (1993).

Received May 1995

Accepted in revised form August 1996 by J. M. Wing