HasCasl: Integrated Higher-Order Specification and Program Development

Lutz Schröder and Till Mossakowski [1]

DFKI-Lab Bremen and Dept. of Comput. Sci., Universität Bremen

Abstract

We lay out the design of HasCasl, a higher order extension of the algebraic specification language Casl that serves both as a wide-spectrum language for the rigorous specification and development of software, in particular but not exclusively in modern functional programming languages, and as an expressive standard language for higher-order logic. Distinctive features of HasCasl include partial higher order functions, higher order subtyping, shallow polymorphism, and an extensive type-class mechanism. Moreover, HasCasl provides dedicated specification support for monad-based functional-imperative programming with generic side effects, including a monad-based generic Hoare logic.

Key words: Algebraic specification, functional programming, type classes, polymorphism, Casl, monads, Hoare logic

Introduction

The rigorous development of software from abstract requirements to executable code calls for wide-spectrum languages that are sufficiently powerful and flexible to support both an expressive specification logic and concepts appearing in advanced programming languages, including modern functional languages such as Haskell [58], but also imperative and object-oriented languages. Here, we discuss the design of such a wide-spectrum language, HasCasl. HasCasl is an extension of the standard algebraic specification language Casl (*Common Algebraic Specification Language*) [6,53] developed by the Common Framework Initiative (CoFI) of IFIP WG 1.3, and as such has been adopted by IFIP WG 1.3. It arguably constitutes 'the' natural higher

[1] Research supported by the DFG project HasCASL (KR 1191/7-1/2) and the BMBF-project FormalSafe (FKZ 01IW07002)

order extension of Casl, and is intended, beyond its purpose as a software specification language, as an expressive standard language for higher order logic. In particular, HasCasl is presently the most expressive language in the logic graph underlying the Bremen heterogeneous tool set Hets [50] and as such serves as a central hub for the interchange of theories between various formalisms in the tool.

The core of HasCasl is a higher order logic of partial functions built on top of Moggi's partial λ-calculus [44]. The semantics and proof theory of this logic have been developed in a companion paper [71]; essentially, one arrives at an intuitionistic partial higher order logic without choice principles (even without unique choice). The full HasCasl logic extends the core logic by subtyping and type-classed based shallow polymorphism, including higher-order type constructors and constructor classes; the semantics of the latter is based on models explicitly incorporating signature extensions [76]. Support for general recursive functions is bootstrapped in the style of HOLCF [64] by specifying a theory of fixed point recursion on complete partial orders. Extensive syntactical sugaring of these concepts yields an executable sublanguage which is in close correspondence with a large subset of Haskell.

Part of the technical difficulties arising in the development of these concepts stem from the above-mentioned lack of unique choice in the core logic; in particular, additional effort is required in the construction of inductive datatypes and in setting up the theory of complete partial orders. We believe that this effort is justified, as making do without unique choice allows keeping the model theory more general and forces the use of simpler constructions. A more detailed discussion of this point is found in Sect. 2.

HasCasl is powerful enough to serve as a framework for the definition of further advanced specification logics. We illustrate this point by developing a generic Hoare logic for reasoning about functional-imperative programs with generic side-effects. Following seminal work by Moggi [46], side effects are encapsulated in functional programming via so-called monads; in particular, this is one of the central concepts of Haskell [80]. Monads model a wide range of computational effects: e.g., stateful computations, non-determinism, exceptions, input, and output can all be viewed as monadic computations, and so can various combinations of these concepts such as non-deterministic stateful computations. Our Hoare logic provides a generic *logical* environment for reasoning about monadic computations. In this way, we generalise the suggestions of [46], which were aimed purely at a state monad with state interpreted as global store. We provide both a generic kernel calculus and specialised calculi that provide additional rules dealing with monad-specific operations such as assignment. We end up with an environment that offers not only a combination of functional and imperative programming (as provided in Haskell), but also a surrounding logic that is rather effortlessly adapted to the specification

of both functional and imperative aspects.

The material is organised as follows. We give a brief introduction to CASL in Sect. 1. We then recall HASCASL's core logic in Sect. 2. Sections 3–5 deal with the syntax and semantics of type class oriented shallow polymorphism, subtyping, and inductive datatypes, respectively. The HOLCF style modelling of general recursion is treated in Sect. 6. In Sect. 7, we discuss the integration of HASCASL into the heterogeneous tool set, in particular its connection with CASL, Isabelle/HOL, and Haskell. Sections 8–11 are concerned with the monad-based generic Hoare calculus. We give an introduction to monad-based functional-imperative programming, and then discuss generic notions of purity and the calculus proper. We illustrate the calculus by means of an extended example, where Dijkstra's non-deterministic implementation of Euclid's algorithm is verified over a generic non-deterministic reference monad.

While we have taken care to keep the presentation self-contained where feasible, this has not always been possible within reasonable space, in particular where the semantic foundations of HASCASL are concerned. Specifically, the following sections assume substantial background knowledge of the reader. Section 2.4 requires familiarity with the categorical semantics of higher order logic, and draws on results from [71]. Readers interested primarily in the HASCASL language design may safely skip this section; where reference to models is made later, exact knowledge of their definition is typically not required. Sections 3.2 and 7 are concerned with institution-theoretic aspects of HASCASL; while we do recall the definitions involved, we refer to [22,23] for more detailed explanations and motivation of the general concepts. Both sections are of interest only to readers with suitable background. Finally, the development of a version of domain theory in HASCASL's internal logic in Sect. 6.1, although technically self-contained, presupposes background knowledge in standard domain theory in terms of motivation. Readers not interested in the domain-theoretic foundations may safely jump to the description of the syntax of executable specifications in Sects. 6.2 and 6.3.

A preliminary version of the HASCASL design has appeared in [73]; the sections on the monad-based generic Hoare logic extend [74].

Related frameworks There are several approaches to tackling the transition from specifications to programs in the literature. Many of them, including Larch [25] and VSE-2 [31], keep the level of specifications in the well-studied realm of first-order logic, while the more problematic features of programming languages are dealt with in intermediate logics (like the dynamic logic of VSE-2) or in programming language specific interface languages (as in Larch). Extended ML [34] avoids such mediating languages by building a higher order specification language on top of a programming language; however, SML's

side effects lead to quite complex interactions with the type system and the logic. A number of such problems is listed in [33]; all these disappear when one moves to a purely functional language like Haskell. Some specific issues mentioned in [33], such as polymorphism over unused type variables and repetition of side-effecting expressions, relate to material discussed in Sects. 3 and 9. A specification logic for Haskell, called P-logic [27], is provided by the Programatica framework [26]. In particular, P-logic resembles our approach in the way polymorphism and recursion are supported; the latter is based on an axiomatic treatment of complete partial orders. P-logic differs from our approach in that it is directly built on top of Haskell (with all its specialities like lazy pattern-matching), while we provide a general-purpose higher-order logic that is both a generalisation of classical higher-order logic *and* can be used as a specification logic for Haskell programs. In particular, HasCasl allows loose requirement specifications that are later refined into design specifications and programs, which is not possible with the P-logic approach. Moreover, HasCasl covers type class based overloading and constructor classes in full generality, whereas P-logic [35] seems to be equipped with specific built-in rules for one particular constructor class, namely monads. Moreover, for monads, P-logic only offers equational reasoning, whereas HasCasl offers a Hoare logic for imperative monad-based programs.

Other approaches such as CaféOBJ [16] and the related tool Maude [12] opt for making the specification language itself directly executable, however at the expense of a reduced expressivity of the logic. VDM [32] and Z [78] are *model-oriented* specification languages, i.e. a specification typically describes one single intended input-output behaviour. By contrast, Casl and HasCasl allow for loose specifications that abstractly describe whole collections of behaviours in order to avoid overspecification in the early phases of the development.

In terms of the logic employed, HasCasl is related in many ways to Isabelle/HOL [54] and Isabelle/HOLCF [64], respectively, with the crucial difference being that HasCasl works with a more flexible logic that does not impose strong reasoning principles such as excluded middle and choice from the outset. Like Coq [14], HasCasl allows adding such principles explicitly as axioms if desired. There is a certain design trade-off here between encapsulating all reasoning principles to be made available in a strong core logic which later can only be extended conservatively or in fact definitionally, as done in Isabelle/HOL, and choosing a weaker core logic but then allowing non-conservative extensions as in Coq or HasCasl – the latter approach is more versatile but also in some sense more dangerous. That said, many constructions which are performed in Isabelle/HOL using unique choice, e.g. datatypes and recursion, are definitional also in HasCasl over comparatively harmless extensions such as sum types and a type of natural numbers, as discussed in Sect. 5.

In fact, Isabelle/HOL presently forms the core of the reasoning support for classical HasCasl. The gap to be bridged here stems mainly from the fact that HasCasl is a specification language aimed at ease of expression, while the logic of Isabelle/HOL is an input language for a proof tool, and as such more austere. Features of HasCasl not directly supported in Isabelle include higher order type constructors and constructor classes (the latter are needed e.g. for modelling side-effects via monads as explained above), subtyping, partial functions, loose generated types and advanced structured specification constructs. Similar comments apply to other higher-order theorem provers such as PVS [56].

Existing dedicated higher-order frameworks for software specification include Spectrum [8] and RAISE [21]. Spectrum is in some ways a precursor of HasCasl, in particular supports higher-order functional programming and offers a type class system (without constructor classes). It is however designed entirely as a language for complete partial orders; consequently, it has a three-valued logic admitting undefined truth values and moreover does not include a proper higher-order *specification* language (i.e. non-continuous functions are included for specification purposes, but higher order mechanisms such as λ-abstraction are limited to continuous functions). The RAISE specification language RSL concentrates on direct support for imperative programming and non-determinism, covered in HasCasl by a monad mechanism. The main differences with HasCasl are that RSL has a three-valued logic and does not support polymorphism.

1 CASL

The specification language Casl (*Common Algebraic Specification Language*) has been designed by CoFI, the international *Common Framework Initiative for Algebraic Specification and Development* [13]. Its features include first-order logic, partial functions, subsorts, sort generation constraints, and structured and architectural specifications. For the definition of the language including a full formal semantics see [53]. An important point here is that the semantics of structured and architectural specifications is institution-independent, i.e. independent of the logic employed for basic specifications. In order to define the envisaged extension of Casl, it is therefore sufficient to define the underlying logic in the form of an *institution* [22], i.e. essentially to fix notions of signature, model, sentence, and satisfaction as done below.

We briefly point out some particularities of the Casl notation that appear in more or less the same form in HasCasl, but refer to [53,6] for a full explanation of the concepts involved. The Casl logic is *multisorted*; the user may declare sorts by means of the keyword **sort**. Sorts appear in the profiles

of operations and predicates. The interpretation of sorts, operations etc. is, by default, *loose*, i.e. a sort may be interpreted by any set and an operation may be interpreted by any map of the given profile, as long as the axioms of the specification are satisfied. The latter, of course, may force an essentially unique interpretation; this holds in particular for axioms implicit in datatype declarations. CASL supports partial operations; the fact that an operation is possibly partial is indicated by a question mark in its profile, i.e. a partial operation f with argument sorts $s_1, \ldots, s_n$ and target sort t is declared in the form

op $f : s_1 \times \cdots \times s_n \rightarrow ?t$

while for a total operation the profile is written in the form $s_1 \times \ldots s_1 \rightarrow t$. There are atomic formulas for definedness: the formula *def* α asserts that the term α is defined. Partial functions are *strict*, i.e. *def* $f(\alpha)$ always implies *def* α. There are two forms of equations between partial terms: a *strong equation* $\alpha_1 = \alpha_2$ asserts that α_1 is defined iff α_2 is defined, and in this case both terms are equal, while an *existential equation* $\alpha_1 \stackrel{e}{=} \alpha_2$ asserts that α_1 and α_2 are defined and equal. Predicate applications impose existential interpretation of their arguments, i.e. they only hold if all arguments are defined. Finally, CASL distinguishes between local and global *variables*: the scope of a global variable, declared using the keyword **var**, is the entire remaining basic specification, while the scope of a local variable, declared using **forall**, is limited to the immediately following list of axioms. Both global and local variables are understood to be universally quantified.

2 The Basic Logic of HASCASL

HASCASL is based on the partial λ-calculus with equality as introduced in [44,45,65]. The model theory of HASCASL uses results of [71] relating the categorical semantics given in [44] with a set-theoretic semantics which is compatible with the existing semantics of CASL.

2.1 The Basic Type System

The (extensible) type system of HASCASL features product types, partial and total function types, and a unit type. Types are built from user-declared basic types introduced by the keyword **type** (for the sake of compatibility with CASL, the keyword **sort** may be used alternatively; moreover, most HASCASL keywords may also be used in their plural forms). E.g. writing

types S, T

declares two basic types S and T. As in CASL, the interpretation of basic types is, by default, loose (Sect. 1). From these basic types and the *unit type Unit*, the *types* are inductively generated by taking *product types* $s_1 \times \cdots \times s_n$ and *partial* and *total function types* $s \to ?t$ and $s \to t$, respectively, with $s \to ?t$ a type of *strict* partial functions (Sect. 1).

A type may be abbreviated by means of a *synonym*, using also the keyword **type**, by writing e.g.

type $Binary := (S \times S) \to T$

The type referred to by a type synonym is called its *expansion*. Although the same keyword is used, synonyms are not basic types. A synonym may be defined only once. Recursive synonym definitions are not allowed.

Terms are formed along with their types; we will introduce the term formation rules informally below. The judgement that a term α has type s is written $\alpha : s$. In fact, term formation depends also on a *context* $\Gamma = (x_1 : t_1, \ldots, x_n : t_n)$ of typed variables $x_i : t_i$ which may include local variables introduced by quantifiers or λ-bindings as well as global variables (introduced by the keyword **var**), so that the complete form of a typing judgement is $\Gamma \rhd \alpha : s$ (term α has type s in context Γ); we will largely omit this aspect here. Terms are built from variables and user-declared *constants*. A constant (or *operation*) f of type s is declared by writing

op $f : s$;

(with the same mechanisms for declaring several constants at once as in CASL). Instead of **op**, the keyword **fun** may be used (**op** and **fun** differ w.r.t. their behaviour under subtyping; see Sect. 4). Since s may be a function type, this provides also a way to declare operations with arguments. As in CASL, constants may be *overloaded*, i.e. a constant $f : s$ is made up of its *name* f and its type s, and the same name may appear in different constants of different types. In fully statically analysed formulas, all constants are explicitly annotated with their types, while they are usually referred to by just their name in the input syntax if the context information suffices for disambiguation. There is a built-in overloaded constant $\stackrel{e}{=}$, called *internal equality*, of type $s \times s \to ?Unit$ for each s, which has a fixed interpretation as equality (due to strictness necessarily existential), exploiting an identification of predicates and partial functions into *Unit* (Sect. 2.2).

The built-in type constructors come with associated term formation rules. Terms of type $s_1 \times \cdots \times s_n$ may be constructed as *tuples* $(\alpha_1, \ldots, \alpha_n)$, where α_i is a term of type s_i for $i = 1, \ldots, n$. The empty tuple () is a term of type *Unit*. Application of a term $\alpha : s \rightarrow ?t$ or $\alpha : s \rightarrow t$ to a term $\beta : s$ is denoted by juxtaposition in the form $(\alpha\,\beta)$ (where application associates to the left, i.e. $((\alpha\,\beta)\,\gamma)$ may be written as $\alpha\,\beta\,\gamma$). Given a term α in a context containing an additional variable $x : s$, the partial function that takes x to the term α is denoted by $\lambda\,x : s \bullet \alpha$. If α is defined for all possible values of x, then $\lambda\,x : s \bullet!\,\alpha$ denotes the corresponding total function of type $s \rightarrow t$; otherwise, the term $\lambda\,x : s \bullet!\,\alpha$ is still well-formed, but fails to denote a defined value — contrastingly, the term $\lambda\,x : s \bullet \alpha$ is always defined.

Equational deduction systems for the partial λ-calculus (sound and complete w.r.t. a semantics discussed in Sect. 2.4) are given in [44,71]. These systems are easily extended (e.g. using the results of [71]) to encompass product types and total function types as featured in HasCasl. The rules include partial versions of (β), (η), and (ξ), where attention has to be paid to definedness issues (e.g. the strong equation $(\lambda x \bullet \alpha)\,\beta = \alpha[\beta/x]$ holds only when β is defined).

Definition 1 A *basic* HasCasl *signature* consists of sets of basic types, type synonyms, and constants, together with a map associating to each type synonym its expansion as defined above. A *morphism* of basic signatures consists of three maps taking constants to constants, type synonyms to type synonyms, and basic types to basic types or type synonyms, respectively; these maps are required to be compatible in the expected sense with types of operations and expansions of type synonyms. (They are *not* required to preserve name equality of constants; cf. however Defn. 38.)

Remark 2 From the user's point of view, the relevance of the notion of signature morphism is mainly that it determines which argument fittings are admissible in instantiations of parametrised specifications [53] and in refinements between specifications [51]. Since the above definition explicitly allows signature morphisms to map basic types to type synonyms, basic types can be instantiated with composite types, albeit at the cost of having to define a type synonym first (allowing basic types to be mapped directly to composite types would strongly increase the number of signature morphisms matching a so-called raw symbol map [53] and thus make symbol maps harder to write and parse). A consequence is that the signature category fails to be cocomplete (while its non-full subcategory consisting of the signature morphisms that map basic types to basic types *is* cocomplete, being essentially the category of models of a Horn theory). However, the pushouts required for instantiating parametrised specifications do exist, which is all that is needed for HasCasl structured specifications.

Another, related use of signature morphisms is for unions of specifications. The semantics of specification unions in [53] is defined in terms of a notion of signature union, that is, a binary partial function on signatures that (if defined) delivers the union of the two signatures, together with the inclusion signature morphisms. Indeed, in [53], also instantiations of parameterised specifications are defined using signature unions that then are also required to be pushouts. Unions of basic HasCasl signatures are defined if they agree on their type synonyms, and in this case, they are formed by uniting their components.

HasCasl provides the following further term forming operations as convenient syntactic sugar:

Let-terms Local bindings are written $let\ x = \alpha\ in\ \beta$, abbreviating $(\lambda\, x \bullet \beta)\, \alpha$. Equivalently, the form $\beta\ where\ x = \alpha$ may be used. Consecutive bindings may be gathered in the form $let\ x_1 = \alpha_1; \dots; x_n = \alpha_n\ in\ \beta$.

Iterated abstraction Consecutive λ-abstractions may be combined in the form $\lambda\, x_1\ x_2 \dots\ x_n \bullet \alpha$, abbreviating $\lambda\, x_1 \bullet!\, \lambda\, x_2 \bullet!\, \dots \lambda\, x_n \bullet \alpha$. Abstraction over an unused variable of type *Unit* may be written in the form $\lambda \bullet \alpha$.

Patterns Variables may be bound within *patterns* in the same way as in (strict) functional programming. In the language introduced so far, this means that variables may be bound to components of tuples; e.g. in the term $let\ (x, y) = \alpha\ in\ \beta$, where $\alpha : s \times t$, x is bound to the first component of α and y to the second component. In the full language, patterns may also contain datatype constructors; see Sect. 5. Patterns may be arbitrarily nested. HasCasl does *not* include built-in projection functions for product types; these can either be user-defined or replaced by pattern matching. In the meta-theory, we do use fst and snd to denote the projections for a binary product type.

Restriction Given terms $\alpha : s$, $\beta : t$, the term $\alpha\ res\ \beta$ abbreviates the term $let\ (x, y) = (\alpha, \beta)\ in\ x$. That is, $\alpha\ res\ \beta$ is defined iff α and β are defined, and in this case equals α. As a special case, β may be a formula (Sect. 2.2); in this case, $\alpha\ res\ \beta$ is defined iff α is defined and β holds.

Typed terms Again as in Casl, terms as well as patterns may be annotated with their intended type in the form $\alpha : s$. This affects only the static analysis of the term in that fewer typing possibilities have to be considered.

Further extensions by abbreviation of both the type system and the term formation rules are discussed in Sect. 2.2 and 2.3.

2.2 The Internal Logic

Partial functions into *Unit* can be regarded as predicates, with definedness corresponding to satisfaction; an example is the equality operator $\stackrel{e}{=}$ mentioned

above. The type $Pred\ s$ is a built-in synonym for $s \to ?\,Unit$ (the CASL notation **pred** $p : t$ is retained as an alternative to **op** $p : Pred\ t$ for the sake of compatibility). The type $Unit \to ?\,Unit$ serves as a type of truth values, with built-in synonym *Logical*. Using the equality operator, one can *define* a full-blown intuitionistic higher order logic, in which formulas are partial terms of type *Unit* and logical operators and quantified formulas are just abbreviations, as follows [70,71] (we use notation from Sect. 2.1).

$$
\begin{aligned}
\top &:= () \\
p \wedge q &:= p\ res\ q \\
p \Rightarrow q &:= ((\lambda \bullet p) \overset{e}{=} \lambda \bullet (p \wedge q)) \\
p \Leftrightarrow q &:= (p \Rightarrow q) \wedge (q \Rightarrow p) \\
\forall y : t \bullet p &:= ((\lambda y : t \bullet\ p) \overset{e}{=} \lambda y : t \bullet\ \top) \\
\bot &:= \forall a : Logical \bullet a\,() \\
\neg p &:= p \Rightarrow \bot \\
p \vee q &:= \forall a : Logical \bullet ((p \Rightarrow a\,()) \wedge (q \Rightarrow a\,())) \Rightarrow a\,() \\
\exists y : t \bullet p &:= \forall a : Logical \bullet (\forall y : t \bullet p \Rightarrow a\,()) \Rightarrow a\,().
\end{aligned}
$$

(In the seemingly asymmetric definition of conjunction, note that for formulas p, q, i.e. partial terms $p, q : Unit$, one actually has a strong equality $\mathsf{fst}(p, q) = \mathsf{snd}(p, q)$.) Atoms are either predicate applications, existential equations $\overset{e}{=}$, definedness assertions $def\ \alpha$ abbreviating $\alpha \overset{e}{=} \alpha$, or strong equations $\alpha = \beta$, defined as

$$(\alpha = \beta) := (def\ \alpha \Rightarrow \alpha \overset{e}{=} \beta) \wedge (def\ \beta \Rightarrow def\ \alpha)$$

Satisfaction of a formula is just definedness of the corresponding term. One thus obtains an extension of the CASL formula syntax. Syntactical differences stem primarily from HASCASL's richer type and term system; in particular, higher order variables can be used in quantifications. The definition of the logic can be written as a HASCASL specification [73]. We refer to this logic as the *internal logic* (the initial design of HASCASL [73] comprised an alternative external logic).

To illustrate some features of the basic logic, we give a pedestrian specification of an abstract while operator as a least fixed point (a more concise specification using a built-in notion of general recursion is given in Fig. 10) in Fig. 1; the specification imports a two-valued type of Booleans contained in the specification SUMS shown further below (Fig. 3).

As mentioned above, the internal logic is intuitionistic: for $p : Logical$, $p \vee \neg p$ is not provable (this fact is unsurprising in view of the fact that the base calculus is essentially just a λ-calculus, and can easily be seen model-theoretically; cf. Sect. 2.4). If desired, the user may impose classical logic by importing the specification

$$
\begin{array}{lll}
\textbf{spec} & \textsc{While} = \textsc{Sums}\ \textbf{with}\ \mathit{Bool}, \mathit{if}\ _\ \mathit{then}\ _\ \mathit{else}\ _ & \textbf{then} \\
 & \textbf{type} \quad s & \\
 & \textbf{ops} \quad \mathit{while} : (s \to \mathit{Bool}) \to (s \to ?\, s) \to s \to ?\, s; & \\
 & \qquad\quad _ \sqsubseteq _ : \mathit{Pred}\ ((s \to ?\, s) \times (s \to ?\, s)) & \\
 & \textbf{vars} \quad b : s \to \mathit{Bool};\ p : s \to ?\, s;\ x : s & \\
 & \qquad\quad \bullet\ _ \sqsubseteq _ = \lambda (q, r) : (s \to ?\, s) \times (s \to ?\, s) \bullet & \\
 & \qquad\qquad \forall y : s \bullet (r\ y)\ \mathit{res}\ (q\ y) = q\ y & \\
 & \qquad\quad \bullet\ \mathit{let}\ F = \lambda q : s \to ?\, s;\ x : s \bullet \mathit{if}\ b\ x\ \mathit{then}\ x\ \mathit{else}\ q\ (p\ x); & \\
 & \qquad\qquad w = \mathit{while}\ b\ p\ \mathit{in} & \\
 & \qquad\qquad\quad F\ w = w\ \wedge \forall q : s \to ?\, s \bullet F\ q = q \Rightarrow w \sqsubseteq q &
\end{array}
$$

Fig. 1. Specification of an abstract while operator

$$
\begin{array}{ll}
\textbf{spec} & \textsc{Classical} = \\
\textbf{var} & p : \mathit{Logical} \\
 & \bullet\ p \vee \neg p
\end{array}
$$

Remark 3 While for many purposes, e.g. inheriting mechanised proof support from Isabelle/HOL, one will often need to work with classical HasCasl (i.e. import the above specification Classical), keeping the base logic intuitionistic opens up a number of possibilities that are not fully available in classical frameworks. E.g., intuitionistic logics offer better facilities for program extraction than classical ones [40,4,42]. Moreover, many useful theories and principles are consistent with intuitionistic but not with classical logics; this includes e.g. set-theoretic parametric polymorphism [63] or an axiom stating that all functions are recursive [60].

A further property of the internal logic is that it distinguishes between *functions*, i.e. inhabitants of function types $a \to ?b$, and *functional relations*, i.e. right-definite predicates on $a \times b$ (see [71] for details). In other words, the internal logic does not in general have a *unique choice* operator ι that, given a formula $x : a \rhd \phi$, returns the unique element $\iota\, x : a \bullet \phi$ of type a satisfying ϕ if a unique such element exists (and is defined iff this is the case). Types a for which such an operator does exist are called *coarse*. Generally, every type of the form $\mathit{Pred}\ a$, including $\mathit{Logical}$, is coarse (as one can put $\iota\, p : \mathit{Pred}\ a \bullet \phi = (\lambda x : a \bullet \forall p : \mathit{Pred}\ a \bullet \phi \Rightarrow p(x))\ \mathit{res}\ \exists!\, p : \mathit{Pred}\ a \bullet \phi$, where $\exists!$ is unique existential quantification, defined by abbreviation as usual), and coarse types are stable under products, function spaces, and subtypes; moreover, every type a has an *underlying* coarse type, the type of singleton subsets of a.

Remark 4 It will become apparent below (in particular in Sect. 5 and 6) that a certain amount of additional effort is required to make standard concepts and constructions work in the absence of unique choice. The motivation justifying this effort is partly to admit certain useful models; this is discussed in

detail in Rem. 11 below. Moreover, a discipline of avoiding unique choice leads to constructions which may be easier to handle in machine proofs than ones containing unique description operators; see e.g. the explicit warning in [54], Sect. 5.10. That said, the user may impose unique choice globally or for selected types: a type a is equipped with unique choice by the specification

op $\quad$ $\mathit{choose} \colon \; (\mathit{Pred}\ a) \;\rightarrow ?\, a$
var $\quad$ $p \colon \mathit{Pred}\ a;\; x : a$
$\quad\quad \bullet\ \mathit{choose}\ p = x \Leftrightarrow (\forall y \colon a \bullet p\ y \Leftrightarrow x = y)$

which be may either imposed on the individual type a or made polymorphic over some class of types (Sect. 3). The terms $\iota\, x : a \bullet \phi$ mentioned above can then be written in the form $\mathit{choose}\ \lambda x : a \bullet\ \phi$.

Remark 5 A much stronger choice operator is Hilbert's ϵ; it corresponds to the Axiom of Choice and implies classicality. It is specified as follows:

op $\quad$ $\mathit{epsilon} \colon \; (\mathit{Pred}\ a) \;\rightarrow a$
var $\quad$ $p \colon \mathit{Pred}\ a;\; x : a$
$\quad\quad \bullet\ (\exists x \colon a \bullet p\ x) \Rightarrow p\ (\mathit{epsilon}\ p)$

The term $\epsilon\, x : a \bullet \phi$ can then be written in the form $\mathit{epsilon}\ \lambda x : a \bullet\ \phi$.

Definition 6 A *basic* HasCasl *theory* is a basic signature together with a set of formulas.

For later use, we fix notions concerning subtypes determined by formulas.

Definition 7 A *generalised type* is a pair $(t.\, \phi)$ consisting of a type t and a predicate $\phi : \mathit{Pred}\ t$ (in the terminology of [71], the generalised types are the objects of the classifying category), to be understood as the subtype of all elements $x : t$ satisfying $\phi\, x$.

Remark 8 Products and (partial or total) function spaces of generalised types can again be described as generalised types [71].

2.3 Non-Strict Functions

In HasCasl, as in the partial λ-calculus, function application is *strict*, i.e. defined values are obtained only from defined arguments. This is in keeping with the semantics of both Casl and ML, but not with the semantics of Haskell, where functions are allowed to leave arguments unevaluated and thus yield defined results on undefined arguments. It is well-known that non-

strict functions may be emulated in a strict setting by moving to function types $Unit \rightarrow ?a$ as argument types. In order to facilitate the specification of programs in non-strict languages, we include non-strict function types in HasCasl as syntactic sugar:

For a type s, $?s$ abbreviates the type $Unit \rightarrow ?s$. Thus, we obtain *non-strict function types* such as $?s \rightarrow ?t$. There are two typing rules for application of non-strict functions and application of strict functions to non-strict values:

- If α is a term of type $?s \rightarrow ?t$ or $?s \rightarrow t$ and β is a term of type s, then $\alpha\,\beta$ is a term of type t, in which β is implicitly replaced by $\lambda \bullet \beta$.
- If α is a term of type $s \rightarrow ?t$ or $s \rightarrow t$ and β is a term of type $?s$, then $\alpha\,\beta$ is a term of type t, in which β is implicitly replaced by $\beta\,()$.

Corresponding generalised rules apply to functions with several arguments. As a simple example, consider the following specification of a non-strict conjunction on the type *Bool* of Booleans (see Fig. 3 below):

op $And : Bool \rightarrow (?Bool) \rightarrow Bool$
var $x : ?Bool$
$\bullet$ $And\ False\ x = False$
$\bullet$ $And\ True\ x = x$
$\bullet$ $And\ False\ True = False$ % **implied**

Here, the last occurrence of $x : ?Bool$ is implicitly replaced by $x()$, while the occurrence of *True* in the last formula (which, as indicated by the Casl annotation %**implied**, is implied by the others) is replaced by $\lambda \bullet True$.

2.4 Model Semantics

We now recall the model-theoretic semantics of HasCasl as developed in [71]. Readers without a background in the semantics of higher-order logics may safely skip this section for most purposes.

The semantics of HasCasl extends the set-theoretic semantics of first order Casl; that is, types are interpreted as sets, and constants are interpreted as elements of these sets. The principal issue is then the interpretation of function types; common options include the following.

- In *standard* models, function types $s \rightarrow t$ and $s \rightarrow ?t$, respectively, are interpreted by the full set of (partial) functions from the interpretation of s to that of t.
- In *extensional Henkin* models [29], function types are interpreted by subsets

of the full set of functions in such a way that all λ-terms can be interpreted; the latter property is called *comprehension*. (In the model theory of the total λ-calculus, similar models are called *λ-models*).
- In *intensional* Henkin models (similar to *λ-algebras*), function types are interpreted by arbitrary sets equipped with an application operation. Comprehension is still required; however, the way λ-terms are interpreted is now part of the structure of the model rather than just an existence axiom. Intensionality is discussed e.g. in [43,44].

The notion chosen for HasCasl is that of intensional Henkin models. Intensional models behave well w.r.t. existence of initial models (unlike extensional models [3]) and, unlike standard models, admit a complete deduction system (completeness for extensional models is at least difficult [44]). Moreover, they are the natural models for *intuitionistic* higher order logic; see Rem. 3 for a brief discussion of the relevance of intuitionistic logic in computer science. (That said, the logic can be specified to be classical by the user if desired; see Sect. 2.2.) We do however introduce variants of HasCasl with extensional and even standard models in order to facilitate the embedding of Casl into HasCasl (Sect. 7).

A peculiarity of the intensional approach is that the definition of model requires an equational deduction system. As indicated in Sect. 2, we assume given an obvious extension of the deduction system presented in [71] with product types and total function types.

Definition 9 An *(intensional Henkin) model* of a given HasCasl signature is an assignment of a set M_s to each type s, in such a way that *Unit* is interpreted as a singleton set and product types are interpreted as cartesian products, together with an assignment of a partial interpretation function

$$M_{s_1} \times \cdots \times M_{s_n} \to ? M_t$$

to each term of type t in context $(x_1 : s_1, \ldots, x_n : s_n)$. These interpretation functions are required to respect deducible equality of terms. Moreover, substitution must be modelled as composition of partial functions, and terms of the form $x_1 : s_1, \ldots, x_n : s_n \rhd x_i : s_i$ must be interpreted by the appropriate product projections. (It follows that tuple terms are interpreted by tupling of functions, and that all total functions that live in the partial function type are represented in the total function type; the latter is proved using total λ-abstraction.)

A *model morphism* between two such models is a family of functions h_s, where s ranges over all types, that satisfies the usual (weak) homomorphism condition w.r.t. all terms. For details, we refer to [71].

By the results of [71], Henkin models are equivalent to models in partial carte-

sian closed categories (pccc's) with equality as defined in [44]; typical examples of pccc's with equality are quasitoposes [82]. Indeed giving a pccc model is often the easiest way to construct a Henkin model: given a pccc **C** with equality, a Henkin model of a signature may be defined by interpreting sorts as objects in **C**, and constants by (global) elements of the arising interpretations of the corresponding types; such a model is called a *model over* **C**.

The intuitionistic character of the internal logic corresponds precisely to the intensional character of models. Indeed, in an *extensional* model, the type $Logical = Unit \rightarrow ?\, Unit$ contains just the two set-theoretic partial functions from *Unit* to itself, corresponding to the truth values *true* and *false*, i.e. the logic becomes classical. In intensional models, on the other hand, *Logical* may have more than one element; in general, the elements of *Logical* form a Heyting algebra. Under the axiom of excluded middle (see Sect. 2.2), this Heyting algebra becomes a Boolean algebra, which however still may have more than two elements; in other words, the axiom of excluded middle does not imply extensionality of function types, including *Logical*. One should keep in mind that intensionality also has a bearing on the definition of satisfaction. E.g., satisfaction of $\forall x : s \bullet \phi$ in a model M is not the same as satisfaction of ϕ by all elements of M_s (rather, it amounts to satisfaction of an equation in an intensional function type; see Sect. 2.2).

The above considerations relate strongly to the well-known observation that toposes are models of intuitionistic set theory (see e.g. [36]). The crucial distinction between general pccc models and topos models is that in models over a topos, all types are coarse, i.e. toposes satisfy unique choice (see Sect. 2.2); conversely, any pccc model with unique choice is a topos [71].

Example 10 The model theory of intuitionistic set theory in toposes, i.e. of HasCasl models with unique choice, is well-established. Well-known toposes are e.g. the category of sets, categories of sheaves and presheaves, Hyland's effective topos (see e.g. [60]), and the category of nominal sets [20], also known as the Schanuel topos.

Simple examples of models with a classical internal logic but without unique choice are obtained as models over set-based quasitoposes, such as the categories of pseudotopological spaces or (reflexive, symmetric) relations [2]. In such models, the coarse types are precisely those that are interpreted as *indiscrete objects*, where an object is indiscrete if all identity maps into it are morphisms. E.g. in the quasitopos of (reflexive, symmetric) relations, a type a is coarse iff its interpretation is a set X equipped with the indiscrete binary relation, i.e. the full relation $X \times X$. Indeed one easily checks that all predicate types, and therefore also their subtypes of singleton sets, carry the indiscrete relation, and therefore the selection map that assigns to each singleton $\{x\}$ its unique element x is relation-preserving iff X also carries the indiscrete

relation. It should be noted that in such models, initial datatypes (Sect. 5) typically fail to be coarse; e.g. the natural numbers object carries the discrete rather than the indiscrete structure (in the case of reflexive binary relations, the discrete structure is the equality relation).

Remark 11 We continue to discuss the motivation for omitting the unique choice axiom in the HasCasl base logic (see Rem. 4). As indicated above, imposing unique choice would amount to limiting the semantics to models over toposes, rather than over the more general quasitoposes. Typical examples of quasitoposes, besides the ones mentioned above, are categories of extensional presheaves, including e.g. the category of reflexive logical relations, and categories of assemblies, both appearing in the context of realisability models [60,66]. In particular, the category of ω-sets is a quasitopos but not a topos; it is embedded as a full subcategory into the effective topos, whose objects however have a much more involved description than ω-sets [60]. Quasitoposes also play a role in the semantics of parametric polymorphism [7]. It thus seems worthwhile to admit quasitopos models.

3 Type Class Polymorphism

On top of the basic HasCasl logic, we now introduce a form of syntax-oriented shallow type class polymorphism. That is, we allow types and operations to depend on type variables, including type constructor variables, and axioms to be universally quantified over types at the outermost level. Type variables are understood syntactically, i.e. as ranging over all types expressible in the signature. Similarly as in Haskell [58] and Isabelle [54], the range of a type variable may be restricted to a given type class, understood as a subset of the syntactical universe of all types. This naive approach, explained in Sect. 3.1, leads to the problem that the institutional satisfaction condition (see Defn. 23 below) fails; this condition essentially requires that satisfaction is invariant under change of notation, i.e. under signature morphisms. We therefore introduce a second level of the semantics, where this defect is repaired by moving to so-called extended models; this is laid out in detail in Sect. 3.2.

3.1 Syntactic Type Class Polymorphism

The class system of HasCasl, like that of Haskell and unlike that of Isabelle, goes beyond simple type classes in that it caters also for type constructors of arbitrary rank. We therefore begin by introducing a suitable kind universe:

Definition 12 From a given set C of *classes*, which includes a class *Type*, the

set K of *kinds* is formed by the grammar

$$K ::= C \mid K \to K.$$

Kinds of the form $Kd_1 \to Kd_2$ are called *constructor kinds.* A kind is called *raw* if it mentions no classes other than *Type* (e.g. $(\mathit{Type} \to \mathit{Type}) \to \mathit{Type})$). A *subclass relation* is a relation $\leq_C$ between classes and kinds, subject to the following condition. Let S denote the congruence (w.r.t. $\to$) generated by $\leq_C$ on the set of kinds; then the S-equivalence class of each class Cl is required to contain a unique raw kind, denoted $\mathsf{raw}(Cl)$ and called the *raw kind* of Cl. It follows that each kind Kd is S-equivalent to a unique raw kind $\mathsf{raw}(Kd)$, obtained by replacing all classes in Kd with their raw kinds.

Intuitively, *Type* is the syntactical universe of all types, constructor kinds are universes of type constructors, and classes are subsets of these universes as prescribed by $\leq_C$.

Example 13 The subclass relation is not required to be well-founded; e.g. $Cl_1 \leq_C Cl_2$, $Cl_2 \leq_C Cl_1$, $Cl_2 \leq_C \mathit{Type}$ is legal, and both Cl_1 and Cl_2 have raw kind *Type*. However, a subclass relation $Cl_1 \leq_C \mathit{Type} \to Cl_2$, $Cl_2 \leq_C \mathit{Type} \to Cl_1$ is illegal, since Cl_1 and Cl_2 do not have a raw kind. Similarly, $Cl_1 \leq_C \mathit{Type} \to Cl_2$, $Cl_2 \leq_K \mathit{Type}$, $Cl_1 \leq_C \mathit{Type}$ is illegal, since *Type* and $\mathit{Type} \to \mathit{Type}$ compete as raw kinds of Cl_1. Finally, a subclass relation $\mathit{Container} \leq_C \mathit{Type} \to \mathit{Type}$, $\mathit{Container} \leq_C \mathit{Ord} \to \mathit{Ord}$, $\mathit{Ord} \leq_C \mathit{Type}$ is legal, with $\mathsf{raw}(\mathit{Container}) = \mathit{Type} \to \mathit{Type}$ (note that according to the subkinding rules introduced below, $\mathit{Ord} \to \mathit{Ord}$ is *not* a subkind of $\mathit{Type} \to \mathit{Type}$).

A class Cl is declared as a subclass of a given kind Kd by writing **class** $Cl < Kd$. Subclasses of constructor kinds are called *constructor classes.* A class may be declared to be a subclass of several kinds, which however all have to be of the same raw kind due to the requirements of Defn. 12. Classes declared without explicit superkinds are subclasses of *Type*. For instance,

classes $\mathit{BoundedOrd} < \mathit{Ord}$;
$\mathit{Functor} < \mathit{Type} \to \mathit{Type}$

declares two classes *BoundedOrd* and *Ord* such that $\mathit{BoundedOrd} \leq_C \mathit{Ord} \leq_C \mathit{Type}$, and a constructor class *Functor*.

The subclass relation, which we denote in the meta-theory by $\leq_C$, is extended to a subkind relation $\leq_K$ on K by the rules given in Fig. 2. (These rules will be extended in Sect. 4; a syntax-directed version of the entire system is given in Appendix C.) By induction over derivations, one shows that $Kd_1 \leq_K Kd_2$ implies that Kd_1 and Kd_2 have the same raw kind.

$$\frac{Cl \leq_C Kd \text{ in } \Sigma}{Cl \leq_K Kd} \qquad \frac{Kd_1 \leq_K Kd_2 \qquad Kd_3 \leq_K Kd_4}{Kd_2 \to Kd_3 \leq_K Kd_1 \to Kd_4}$$

$$\frac{}{Kd \leq_K Kd} \qquad \frac{Kd_1 \leq_K Kd_2 \qquad Kd_2 \leq_K Kd_3}{Kd_1 \leq_K Kd_3}.$$

Fig. 2. Subkinding rules

Type variables are declared along with their kind either by means of the keyword **var** or in local universal quantifications at the outermost level (we will use mostly the former style here). Type variables may then be used in place of types or type constructors, thus making the entity (type, operation, or axiom) where they appear polymorphic over the given kind. For a type, this means that one obtains a *type constructor*; i.e by writing

type $\quad t\ a_1 \ldots a_n : Kd$

where $a_1, \ldots, a_n$ are type variables of kinds $Kd_1, \ldots, Kd_n$, respectively, one declares a type constructor t of kind $Kd_1 \to \cdots \to Kd_n \to Kd$ (in particular of kind Kd when $n = 0$). Polymorphic operations are assigned *type schemes* in the usual sense of shallow polymorphism, i.e. types that are quantified over type variables at the outermost level — HasCasl does *not* admit nested type quantification as in System F. Finally, polymorphic axioms are implicitly universally quantified over their free type variables.

A simple example is a specification of sum types, shown in Fig. 3. It declares a type constructor *Sum* of kind *Type* → *Type* → *Type*, as well as polymorphic operations *inl*, *inr*, and *sumcase*, governed by polymorphic axioms. We immediately obtain the type *Bool* of Booleans as *Sum Unit Unit*, with the if-then-else construct arising as a special case of the case construct. Moreover, the specification declares a universal undefined constant *bot*, which has the effect of making the generalised type $(Unit, \lambda \bullet false)$ an initial object in the model category. Using *bot*, one can define partial extraction functions *outl*, *outr* for sums as shown in the specification. Actually, this extension of the specification is *definitional* (indicated by the annotation %**def**), meaning that each model of the smaller specification can uniquely be extended to a model of the extended specification.

Instances of polymorphic operations may be formed explicitly using square brackets; e.g. given basic types S and T, we have an instance $sumcase[S, S, T]$ of type $(S \to ?T) \to (S \to ?T) \to (Sum\ S\ S) \to ?T$ (the order of the type arguments is determined by the order of declaration of type variables). However, types of instances may also be automatically inferred, so that instances can be referred to by just the name of operation, as done in the above specification. Note that instances may involve types containing type variables, as in

spec SUMS =
  vars $a, b, c\colon Type$
  type $Sum\ a\ b$
  ops $inl\colon a \to Sum\ a\ b;$
   $inr\colon b \to Sum\ a\ b;$
   $sumcase\colon (a \to? c) \to (b \to? c) \to Sum\ a\ b \to? c$
   $bot\colon ?a$
  vars $f\colon a \to? c;\ g\colon b \to? c;\ h\colon Sum\ a\ b \to? c$
   $\bullet\ h = sumcase\ f\ g \Leftrightarrow$
    $(\forall x\colon a;\ y\colon b \bullet h\ (inl\ x) = f\ x \wedge h\ (inr\ y) = g\ y)$
   $\bullet\ \neg def\ \ bot\colon a$
   $\bullet\ sumcase\ inl\ inr = \lambda z\colon Sum\ a\ b \bullet z$ % **implied**
  then %**def**
  vars $a, b\colon Type$
  ops $outl\colon Sum\ a\ b \to? a;$
   $outr\colon Sum\ a\ b \to? b$
   $\bullet\ outl = sumcase\ (\lambda x\colon a \bullet x)\ (\lambda y\colon b \bullet bot)$
   $\bullet\ outr = sumcase\ (\lambda x\colon a \bullet bot)\ (\lambda y\colon b \bullet y)$
  type $Bool := Sum\ Unit\ Unit$
  var $p\colon Bool;\ x, w\colon a$
  ops $True, False\colon Bool;$
   $if_\ then_\ else_\ \colon Bool \times a \times a \to a$
   $\bullet\ True = inl\ ()$
   $\bullet\ False = inr\ ()$
   $\bullet\ if\ p\ then\ x\ else\ w = sumcase\ (\lambda \bullet x)\ (\lambda \bullet w)\ p$

Fig. 3. A specification of sum types and an implicit initial object

the %**implied** formula in Fig. 3, where the instance of *sumcase* is for a, b, $Sum\ a\ b$.

Remark 14 Since polymorphic overloading is permitted, explicit instantiations as explained above may be ambiguous in case the given arguments fit more than one polymorphic profile. In this case, the use of the operation is disambiguated by its own expected type (cf. Rem. 21); if necessary, explicit type annotations must be given. Partial explicit instantiation, e.g. by writing $sumcase[S, S]$ in the above example, is not allowed.

Product and function types are regarded as applications of built-in type constructors $\times$, $\to$, $\to?$ of kind $Type \to Type \to Type$, and the unit type as a type 'constructor' of kind $Type$; all other type constructors are called *user-declared*. Type constructors are distinguished from parametrised type synonyms; e.g.

var $a\colon Type$
type $DList\ a := List\ (List\ a)$

defines a parametrised type synonym *DList* of kind $Type \to Type$. In closed form, the expansion of *DList* is the *pseudotype*

$$\lambda\, a : Type \bullet List\ (List\ a).$$

In order not to overtax the user with yet another kind of λ-abstraction, such λ-types are not included directly in the syntax of HasCasl. They do however play a role for the range of type constructor variables; see Sect. 3.2.

Kinds are essentially purely syntactic entities: a kind *Kd* does not have a semantics beyond the set of *instances* derivable for it, i.e. the set of types or pseudotypes of kind *Kd*. The instances are governed by two mechanisms: the kinds assigned to type constructors, and the subclass relation. E.g.

var $a : Ord$
type $List\ a, Nat : Ord$

declares the type *Nat* to be of class *Ord* and the type constructor *List* to be of kind $Ord \to Ord$. The latter statement means that *List t* is of class *Ord* whenever *t* is of class *Ord*. A type constructor may be given any number of kinds, which however are required to have the same raw kind.

To a class, one can attach both operations and axioms by using appropriate type variables. E.g. the standard operations and axioms for the two classes of orders declared above are specified by

vars $a : Ord;\ b : BoundedOrd$
ops $_ \leq _ : Pred\ (a * a);$
$bottom, top : b$
vars $x, y, z : a;\ v : b$
$\bullet\ x \leq x$
$\bullet\ x \leq y \wedge y \leq z \Rightarrow x \leq z$
$\bullet\ x \leq y \wedge y \leq x \Rightarrow x = y$
$\bullet\ bottom \leq v$
$\bullet\ v \leq top$

Then e.g. the comparison operator $\leq$ has instances only for types of class *Ord* — e.g. given the type declarations above, for *Nat*, *List Nat*, *List* (*List Nat*) etc. — and also the order axioms hold only at these types. Note that, given the declarations so far, the class *BoundedOrd* has no instances at all, so that no type has operations *bottom* and *top* (however, types of class *BoundedOrd* may be declared later); more about this point in Sect. 3.2.

Unlike in Isabelle and Haskell, axioms and operations, respectively, do not as such form part of the definition of a class in HasCasl. The effect of Isabelle's

axiomatic type classes can however be emulated as follows. To a class declaration, operations and axioms may be attached in a block marked by curly brackets $\{\ldots\}$, thus declaring the *interface* of the class. E.g. the above declaration of operations and axioms for partial orders could be tied to the class declaration by writing

class $\quad$ *Ord* **{**
vars $\quad a\colon Ord;\ x : a$
op $\quad _\leq_ \ : Pred\ (a \times a)$
$\quad\quad \bullet\ x \leq x \ldots$ **}**

Then, declarations of subclasses of *Ord* and declarations of type constructors with result class *Ord* can be marked with the keyword **instance**, thus generating a proof obligation similar to Casl's %**implied** annotation which states that the interface axioms of the class follow from the axioms for the type or the subclass, respectively, together with the axioms of the local environment including the class axioms for type arguments. Here, 'follow' refers to the notion of semantic consequence on the second level of the semantics as introduced in Sect. 3.2 below; this notion of consequence is however just the intuitively expected one (and also the one employed in proof systems such as Isabelle). E.g., given the class interface for *Ord*, we can declare a generic instance for product types by

vars $\quad a, b\colon Ord$
type instance $(a \times b)\colon Ord$
var $\quad x, y\colon a;\ v, w : b$
$\quad\quad \bullet\ (x, v) \leq (y, w) \Leftrightarrow x \leq y \wedge v \leq w$

which gives rise to the proof obligation that reflexivity, transitivity, and antisymmetry of $\leq$ on $a \times b$ follow from the corresponding laws on a and b and the definition of $\leq$ on $a \times b$. Similarly, proof obligations may be generated for subclasses. E.g.,

class instance *DiscreteOrd* $<$ *Ord*
vars $\quad a\colon DiscreteOrd;\ x, y\colon a$
$\quad\quad \bullet\ x \leq y \Leftrightarrow x = y$

expresses that the order axioms follow from the definition of $\leq$ for discrete orders.

As indicated above, HasCasl supports polymorphism over higher kinds. As an example involving constructor classes and type constructors of higher rank, a specification of the classes of monads and monad transformers as used in the

$$
\begin{array}{lll}
\textbf{spec} & \textsc{Functor} \ = & \\
& \textbf{class} & \mathit{Functor} < \mathit{Type} \to \mathit{Type}\ \{ \\
& \textbf{vars} & F\colon \mathit{Functor};\ a, b, c\colon \mathit{Type} \\
& \textbf{op} & \mathit{map}\colon (a \to b) \to F\ a \to F\ b \\
& \textbf{vars} & x\colon F\ a;\ f\colon a \to b;\ g\colon b \to c \\
& & \bullet\ \mathit{map}\ (\lambda y\colon a \bullet!\ y)\ x = x \\
& & \bullet\ \mathit{map}\ (\lambda y\colon a \bullet!\ g\ (f\ y))\ x = (\mathit{map}\ g)\ (\mathit{map}\ f\ x) \\
& & \}
\end{array}
$$

Fig. 4. The constructor class of functors

Haskell libraries [58] is shown in Fig. 5. The specification of monads slightly modifies the usual definition (see Sect. 8); the axiomatisation of monad transformers follows [38]. The class instance mechanism is illustrated by declaring

$$
\begin{array}{lll}
\textbf{spec} & \textsc{Monad} = \textsc{Functor}\ \textbf{then} & \\
& \textbf{class} & \mathit{Monad} < \mathit{Type} \to \mathit{Type}\ \{ \\
& \textbf{vars} & m\colon \mathit{Monad};\ a, b, c\colon \mathit{Type} \\
& \textbf{ops} & _ \gg= _\colon m\,a \times (a \to ?\, m\,b) \to ?\, m\,b; \\
& & _ \gg= _\colon m\,a \times (a \to m\,b) \to m\,b; \\
& & \mathit{ret}\colon a \to m\,a \\
& \textbf{vars} & x, y\colon a;\ p\colon m\,a;\ q\colon a \to ?\, m\,b;\ r\colon b \to ?\, m\,c;\ f\colon a \to ?\, b \\
& & \bullet\ \mathit{def}\ \ (q\ x) \Rightarrow ((\mathit{ret}\ x) \gg= q) = q\ x \\
& & \bullet\ p \gg= (\lambda x\colon a \bullet \mathit{ret}\ (f\ x) \gg= r) = \\
& & \qquad\qquad\qquad p \gg= (\lambda x\colon a \bullet r\ (f\ x)) \\
& & \bullet\ (p \gg= \mathit{ret}) = p \\
& & \bullet\ ((p \gg= q) \gg= r) = (p \gg= (\lambda x\colon a \bullet q\ x \gg= r)) \\
& & \bullet\ ((\mathit{ret}\ x)\colon m\ a) = \mathit{ret}\ y \Rightarrow x = y \\
& & \} \\
& \multicolumn{2}{l}{\textbf{class instance}\ \mathit{Monad} < \mathit{Functor}} \\
& \textbf{vars} & m\colon \mathit{Monad};\ a, b\colon \mathit{Type};\ f\colon a \to b;\ x\colon m\ a; \\
& & \bullet\ \mathit{map}\ f\ x = x \gg= \lambda y\colon a \bullet \mathit{ret}\ (f\ y) \\
& & \\
\textbf{spec} & \textsc{MonadTransformer} = \textsc{Monad}\ \textbf{then} & \\
& \textbf{class} & \mathit{MonadT} < \mathit{Monad} \to \mathit{Monad}\ \{ \\
& \textbf{vars} & t\colon \mathit{MonadT};\ m\colon \mathit{Monad};\ a\colon \mathit{Type}; \\
& \textbf{op} & \mathit{lift}\colon m\ a \to t\ m\ a \\
& \textbf{vars} & x\colon a;\ p\colon m\ a;\ b\colon \mathit{Type};\ q\colon a \to ?\ m\ b \\
& & \bullet\ \mathit{lift}\ (\mathit{ret}\ x) = (\mathit{ret}\ x)\colon t\ m\ a \\
& & \bullet\ \mathit{lift}\ (p \gg= q) = (\mathit{lift}\ p) \gg= \lambda y\colon a \bullet ((\mathit{lift}\ (q\ y))\colon t\ m\ a) \\
& & \}
\end{array}
$$

Fig. 5. The classes of monads and monad transformers

Monad to be an instance of the constructor class *Functor*, whose definition is shown in Fig. 4.

Remark 15 One should note that every declaration of a polymorphic type, operation, or axiom is *separately* implicitly quantified over all presently declared free type variables. E.g. in Fig. 4, the operation *map* is immediately made polymorphic over a and b, so that instances of it can be used in the following axioms at types other than just a and b – such as $map[a, a]$ in the first axiom, and $map[a, c]$, $map[b, c]$ in the second axiom.

Formally, the syntax of polymorphism is captured as follows.

Definition 16 A *polymorphic* HasCasl *signature* Σ consists of a set of classes with a subclass relation according to Defn. 12, kinded type constructors, type synonyms with given expansions, and typed polymorphic operations as described above (basic types are regarded as type constructors without type arguments, similarly for non-polymorphic operations). Kinds of type constructors are subject to the above-mentioned restriction to coincident raw kinds.

A *morphism* of polymorphic signatures consists of maps taking classes to classes, operations to operations, type synonyms to type synonyms, and type constructors to type constructors or type synonyms, respectively. These maps are required to preserve the subclass relation and the kinding of type constructors in the sense that declared subclass relations and kind assignments for type constructors are mapped to derivable subkinding and kinding judgements, respectively, in the target signature, and to be compatible in the expected sense with expansions of type synonyms and types of constants. Moreover, morphisms must preserve polymorphic overloading: if the source signature contains constants $f : t$ and $f : s$, where t and s are unifiable polymorphic profiles, then the images of $f : s$ and $f : t$ must have the same name in the target signature.

Instances of classes are determined by a kinding system for pseudotypes. Kinding takes place in a *type context* Θ of type variables with assigned kinds according to the (syntax-directed) rules shown in Fig. 6 (to be extended in Sect. 4). A judgement of the form $\Theta \rhd t : Kd$ is to be read 't is a type constructor of kind Kd in context Θ'. In the rules, s and t range over pseudotypes and F over basic type constructors; the premise $F : Kd_1$ in Σ means that Kd_1 is an explicitly assigned kind of F. The introduction rule for type constructors applies also to the built-in type constructors $\times$, $\to$, $\to?$, *Unit*. A *type* in the polymorphic language is a pseudotype of kind *Type*.

By induction over the type structure, one shows that all kinds derivable for a pseudotype t are of the same raw kind, the *raw kind* of t.

Remark 17 Using the subkinding rules of Appendix C, one shows by induction over the type structure that the kinds derivable for a pseudotype are upwards closed w.r.t the subkind relation (which is why we can require an exact fit in the application rule). Moreover, kinding obeys a substitution lemma

$$\frac{F : Kd_1 \text{ in } \Sigma \quad Kd_1 \leq_K Kd_2}{\Theta \rhd F : Kd_2} \qquad \frac{a : Kd_1 \text{ in } \Theta \quad Kd_1 \leq_K Kd_2}{\Theta \rhd a : Kd_2}$$

$$\frac{\Theta \rhd t : Kd_1 \quad \Theta \rhd s : Kd_1 \to Kd_2}{\Theta \rhd s\, t : Kd_2} \qquad \frac{\Theta, a : Kd_1 \rhd t : Kd_2 \quad Kd_3 \leq_K Kd_1}{\Theta \rhd \lambda\, a : Kd_1 \bullet t : Kd_3 \to Kd_2}$$

Fig. 6. Kinding rules for type constructors

and hence is invariant under β-equality (but not under η-equality, which is therefore not imposed on type constructors).

The *first level* of the semantics of polymorphism is defined by reduction to the basic language of Sect. 2 as follows.

Definition 18 The *generalised pseudotypes* of a polymorphic HasCasl signature Σ are formed by the rules of Fig. 6 and an additional rule stating that generalised types (Defn. 7) are generalised pseudotypes of kind *Type*. A *type instance* is a closed generalised pseudotype. The point here is that generalised types may appear as arguments of *user-declared* type constructors — by Rem. 8, we may assume that type instances do not contain applications of the built-in type constructors $\times$, $\to$, $\to ?$ to generalised types. A *loose type* is a type instance of kind *Type* which is an application of a user-declared type constructor.

From Σ, we construct a basic theory $\mathsf{B}(\Sigma)$. The sorts of $\mathsf{B}(\Sigma)$ are the loose types of Σ. Operations are translated as follows. Let f have the type scheme $\forall a_1 : Kd_1, \dots, a_n : Kd_n \bullet t$, and let $s_1 : Kd_1, \dots, s_n : Kd_n$ be type instances. By Rem. 8, the type instance $t[s_1/a_1, \dots, s_n/a_n]$ can be interpreted as a generalised type $v = (u_{s_1,\dots,s_n} . \, \phi_{s_1,\dots,s_n})$ in $\mathsf{B}(\Sigma)$. The operation f is then translated into a collection of operations $f_v : u_{s_1,\dots,s_n}$, where $s_1, \dots, s_n$ range over all type instances, together with axioms $\phi_{s_1,\dots,s_n}\, f_v$. Given this construction, it is clear that every morphism $\sigma : \Sigma_1 \to \Sigma_2$ of polymorphic signatures induces a theory morphism $\mathsf{B}(\sigma) : \mathsf{B}(\Sigma_1) \to \mathsf{B}(\Sigma_2)$.

The *first level* of the semantics is then given as follows. The *models* and *model morphisms* of Σ are those of $\mathsf{B}(\Sigma)$. Let $\forall a_1 : Kd_1, \dots, a_n : Kd_n \bullet \phi$ be a Σ-formula, where ϕ does not contain further quantification over types. Given type instances $s_1, \dots, s_n$, the formula $\phi[s_1/a_1, \dots, s_n/a_n]$ is translated into a formula $\phi^{s_1,\dots,s_n}$ over $\mathsf{B}(\Sigma)$ by replacing polymorphic operations with the appropriate instances and by eliminating generalised types from quantifiers, i.e. by replacing $\forall x : (t . \, \psi) \bullet \chi$ with $\forall x : t \bullet \psi\, x \Rightarrow \chi$, similarly for $\exists$. A Σ-model M *satisfies* $\forall a_1 : Kd_1, \dots, a_n : Kd_n \bullet \phi$ if M, regarded as a model of $\mathsf{B}(\Sigma)$, satisfies $\phi^{s_1,\dots,s_n}$ for all type instances $s_1, \dots, s_n$.

Remark 19 The polymorphism introduced above is essentially shallow polymorphism. The discourse in [15] may create the impression that the combination of shallow polymorphism and higher order logic is inconsistent. However, this is not the case: as demonstrated above, shallow polymorphism can be coded out by just replacing polymorphic operations and axioms by all their instances. The derivation of Girard's paradox in [15], Sect. 5, is based on the assumption that terms of the language are identified up to untyped β-equality in the absence of type annotations; such an equality is obviously unsound w.r.t. the usual notions of model, and the paradox shows that a language with such an equality is inconsistent. When, as in the usual versions of shallow polymorphism, instantiations of polymorphic constants are internally annotated with their types, the contradiction disappears (i.e. its derivation just produces a type error).

Remark 20 In Sect. 2.1, the semantics of let-terms is given using λ-abstraction. This base definition precludes ML-style polymorphism, that is, polymorphic type variables that are universally quantified locally to the let-term. However, this type of non-recursive ML-style polymorphism can be coded out by using a separate λ-variable for each instance that is used in the body of the let-term. The same works also for letrec-terms that do not use true polymorphic recursion, which becomes relevant in program blocks (see Sect. 6.3).

Remark 21 Notice that in the above definition, instances of operations are distinguished by their own types, not by the involved type arguments. This means in particular that polymorphic operations declared with identical names but different profiles agree where their profiles overlap. This is to been seen independently of the fact that for the sake of syntactic convenience, explicit instantiation of polymorphic operations is via their type arguments, which are usually simpler than the type of the operation itself.

For instance, one may write

classes $\mathit{Ord}, \mathit{Num}$
vars $a\colon \mathit{Ord};\ b\colon \mathit{Num}$
ops $\mathit{min}\colon a \times a \to ?\, a;$
$\mathit{min}\colon b \times b \to ?\, b$

thus giving the function min the two polymorphic profiles $\forall a\colon \mathit{Ord} \bullet a \times a \to ?a$ and $\forall a\colon \mathit{Num} \bullet a \times a \to ?a$. By Defn. 18, instances for these two profiles at types belonging to both Ord and Num agree. Similarly, overlapping instances of unifiable profiles agree. E.g. one might sensibly first define a polymorphic extension ordering on partial function spaces, and then declare this ordering to be an instance of the class Ord:

vars $a, b\colon Type$
op $_ \leq _\colon Pred\ ((a \rightarrow ?\, b) \times (a \rightarrow ?\, b))$
. . . %% *Definition of the extension ordering*
type instance $a \rightarrow ?\, b\colon Ord$

One then has two explicit profiles for $\leq$, namely $\forall a\colon Ord \bullet Pred\ (a \times a)$ and $\forall a, b\colon Type \bullet Pred\ ((a \rightarrow ?b) \times (a \rightarrow ?b))$, and the instances of the two operations at the types $Pred\ ((a \rightarrow ?b) \times (a \rightarrow ?b))$ are identical.

Remark 22 We recall that polymorphic definitions may introduce inconsistencies if the entity to be defined depends on fewer type variables than the defining entity (see also [33]). E.g. extending the specification

var $a\colon Type$
type $Flag\ a$
ops $mkf\colon Logical \rightarrow Flag\ a;\ getl\colon Flag\ a \rightarrow Logical$
vars $x\colon Flag\ a;\ b\colon Logical$
$\bullet\ mkf\ (getl\ x) = x$
$\bullet\ getl\ (mkf\ b\colon Flag\ a) = b$
op $sg\colon Flag\ a = mkf\ (\forall x, y\colon a \bullet x = y)$

with the 'definition'

op $c\colon Logical = getl\ (sg\colon Flag\ a)$

is obviously inconsistent. For this reason, such definitions are excluded e.g. as `constdefs` in Isabelle, although the same formulas are admissible as `axioms`. We allow them in HasCasl, in keeping with a general philosophy of analysing only a posteriori which axioms are definitions (and the above axiom would not be classified as a definition by such an analysis).

3.2 The Extended Model Semantics

As mentioned above, the first level of the semantics of polymorphic HasCasl as defined in the preceding section fails to constitute an institution. We now briefly recall the notion of institution, and discuss the failure of the satisfaction condition at the first level of the semantics. We then go on to define a second level of the semantics which does constitute an institution, making use of a general institution theoretic construction introduced in [76].

Definition 23 [22] An *institution* consists of

- a category of *signatures* and *signature morphisms*;
- a contravariant *model functor* assigning to each signature Σ a category $\mathbf{Mod}(\Sigma)$ of *models* and *model morphisms* and to each signature morphism $\sigma : \Sigma_1 \to \Sigma_2$ a *reduct functor* $\mathbf{Mod}(\Sigma_2) \to \mathbf{Mod}(\Sigma_1)$, whose action on models is denoted by $M \mapsto M|_\sigma$, where $M|_\sigma$ is called the *σ-reduct* of M;
- a covariant *sentence functor* assigning to each signature Σ a set $\mathbf{Sen}(\Sigma)$ of sentences and to each signature morphism $\sigma : \Sigma_1 \to \Sigma_2$ a *translation* $\mathbf{Sen}(\Sigma_1) \to \mathbf{Sen}(\Sigma_2)$, whose action is denoted by $\phi \mapsto \sigma\phi$; and
- for each signature Σ, a satisfaction relation $\models$ on $\mathbf{Mod}(\Sigma) \times \mathbf{Sen}(\Sigma)$

such that for each signature morphism $\sigma : \Sigma_1 \to \Sigma_2$, the *satisfaction condition*

$$M \models \sigma\varphi \iff M|_\sigma \models \varphi$$

holds for all $M \in \mathbf{Mod}(\Sigma_2)$ and all $\phi \in \mathbf{Sen}(\Sigma_1)$.

For HasCasl, we have assembled most of the data required by this definition in the preceding section, with the exception of model reduction and sentence translation; these data are completed as follows.

Definition 24 Recall that a signature morphism $\sigma : \Sigma_1 \to \Sigma_2$ of polymorphic HasCasl signatures (Defn. 18) induces a signature morphism $\mathsf{B}(\sigma) : \mathsf{B}(\Sigma_1) \to \mathsf{B}(\Sigma_2)$ between the associated basic HasCasl signatures. A basic signature morphism induces a reduct functor in the usual way (i.e. $M|_\sigma$ interprets symbols by the interpretations of their σ-translations in M). The model classes of the Σ_i are, by definition, those of the $\mathsf{B}(\Sigma_i)$; the reduct functor for σ is defined to be that of $\mathsf{B}(\sigma)$. The translation map for σ works in the obvious way by replacing all symbols in a formula by their images under σ, which may involve replacing type constructors by type synonyms.

Given these definitions, it is clear that the satisfaction condition fails for the first level of the semantics: e.g., a signature morphism σ may be an inclusion $\Sigma_1 \hookrightarrow \Sigma_2$ into a signature with more types, and thus given a Σ_2-model M, a Σ_1-formula of the form, say, $\forall a : \mathit{Type} \bullet \phi$ may hold for $M|_\sigma$, i.e. for the type instances in Σ_1, but fail to hold for M, i.e. for the additional type instances in Σ_2. As an extreme example, which also illustrates that the notion of model at the first level fails to capture the full intuitive meaning of polymorphic specifications, consider the specification of monads in Fig. 5. This specification introduces only a class, but no types of that class. Therefore, the given axioms for monads do not have any instances, i.e. the specification is, at the first level, model-theoretically vacuous, which is certainly not the intended meaning. A model of an extension of the signature of Monad where instances of the class *Monad* are declared may very well violate the monad axioms, while its reduct to the signature of Monad will still trivially satisfy them.

All this is remedied at the second level of the semantics, where models are defined to be first-level models of 'future' extensions of the present signature. This is an instantiation of a generic construction presented in [76]. The formalisation of "future extension" also requests for a way of extending models to the extended signatures. This is done using the notion of derived signature morphisms. The intuition is that a derived signature morphism may map symbols of the source signature not only to symbols, but also to more complex terms in the target signature.

Definition 25 A *derived signature morphism* into a signature Σ is a (standard) signature morphism into an extension $\bar{\Sigma}$ of Σ defined as follows.

- The classes of $\bar{\Sigma}$ are the classes of Σ, and additionally all *homogeneous* sets A of type instances in Σ, i.e. sets A such that all elements of A have the same raw kind.
- The type constructors of $\bar{\Sigma}$ are the type instances of Σ. The kinds assigned to such a type constructor are those derivable according to the rules of Fig. 6 from the kind assignments of Σ and additional kind assignments $t : A$ for every homogeneous set A of type instances and every $t \in A$. Note that the closed types of $\bar{\Sigma}$ may be identified with the type instances of Σ by collapsing layers of type formation.
- For a class Cl and a kind Kd of $\bar{\Sigma}$, $Cl < Kd$ iff all instances of Cl are also instances of Kd.
- The operation constants in $\bar{\Sigma}$ of (closed) type t are the closed terms of the type instance t in Σ. These are determined in a combined calculus involving typing and logical deduction which includes the standard typing and deduction rules and additional rules for generalised types $(t.\, \phi)$,

$$\frac{\alpha : t \quad \phi\ \alpha}{\alpha : (t.\phi)} \quad \text{and} \quad \frac{\alpha : (t.\phi)}{\alpha : t \quad \phi\ \alpha}.$$

Remark 26 The categorically-minded reader will recognise derived signature morphisms as Kleisli morphisms for a corresponding monad on the category of signatures. Note that the classes of Σ have to be retained explicitly in $\bar{\Sigma}$ as they cannot be identified qua syntactic entities with sets of types, while e.g. operation constants can just be regarded as particular terms. Crucially, the identities in the category of derived signature morphisms map every class to itself rather than to the set of its instances.

Definition 27 Any (first-level) Σ-model can be uniquely extended to a $\bar{\Sigma}$-model by just interpreting the closed terms serving as operation constants via term evaluation. Note that the additional classes and type constructors in $\bar{\Sigma}$ do not lead to additional type instances.

The *reduct* of a model against a derived signature morphism into a signature Σ is obtained by first extending to the model to $\bar{\Sigma}$ and then taking ordinary

reduct. In order to define a generalised sentence translation, the notion of sentence has to be slightly extended to include universal quantification over type variables ranging over given sets of type instances, with the obvious semantics; given this extended syntax, the definition of translation is straightforward.

Remark 28 Derived signature morphisms are of independent interest, as they can form the basis for constructor implementations in the sense of [68,67]. Specifically, taking reducts along a derived signature morphism is a constructor that may be used to refine a more complex specification to a simpler one; the terms and type sets involved in the derived signature morphism then determine how to extend a model of the simpler specification to a model of the more complex one.

Definition 29 The *second level* of the semantics of polymorphic HasCasl is defined as follows. The notions of signature and sentence remain unchanged. An *extended model* of a signature Σ_1 is a pair (N, σ), where $\sigma : \Sigma_1 \to \Sigma_2$ is a derived signature morphism in the sense defined above and N is a (first-level) Σ_2-model. The *reduct* $(N, \sigma)|_\tau$ of (N, σ) along a signature morphism τ is $(N, \sigma \circ \tau)$. The extended model (N, σ) *satisfies* a sentence ϕ if

$$N \models \sigma\phi$$

at the first level.

By the results of [76], we have

Theorem 30 *The second level of the semantics of polymorphic* HasCasl *constitutes an institution.*

Moreover, instead of the pathologies indicated above, we obtain

Theorem 31 *On the second level of the semantics, the* Σ*-sentence* $\forall b_1 : Kd_{21}; \dots; b_m : Kd_{2m} \bullet \psi$ *is a semantic consequence of* $\forall a_1 : Kd_{11}; \dots; a_n : Kd_{1n} \bullet \phi$*, where the* Kd_{ij} *are kinds, and* ϕ *and* ψ *are formulas not containing further quantification over type variables, iff*

$$(\forall a_1 : Kd_{11}; \dots; a_n : Kd_{1n} \bullet \ \phi) \models \psi$$

on the first level, equivalently on the second level, in the signature obtained from Σ *by adding type constructors* $b_i : Kd_{2i}$*,* $i = 1, \dots, m$*.*

This is precisely the notion of semantic consequence one would intuitively expect, and also the basis for polymorphic proofs as conducted e.g. in Isabelle [54]. A consequence of the theorem is that the sound and complete proof systems for the partial λ-calculus presented in [44,71] lead to sound and complete proof systems for the second level of the semantics.

A further issue in this context are *model-expansive* or, in CASL terminology, (model-theoretically) *conservative* extensions.

Definition 32 A *theory* in a given institution is a pair $T = (\Sigma, \Phi)$ consisting of a signature Σ and a set Φ of Σ-sentences. A *model* of T is a Σ-model M such that $M \models \Phi$. A signature morphism $\sigma : \Sigma_1 \to \Sigma_2$ is a *theory morphism* $(\Sigma_1, \Phi_1) \to (\Sigma_2, \Phi_2)$ if

$$\Phi_2 \models \sigma\Phi_1,$$

where $\models$ denotes logical consequence. We say that σ is *(model-theoretically) conservative* if every model M of (Σ_1, Φ_1) has a (Σ_2, Φ_2)-extension, i.e. a model M' of (Σ_2, Φ_2) such that $M'|_\sigma = M$.

It is easy to see that conservative theory morphisms at the second level are sections (i.e. have left inverses) as derived signature morphisms; conversely, theory morphisms which are sections as derived theory morphisms are conservative [76]. Informally, this means that extensions by syntactic definition are conservative, where thanks to the use of derived signature morphisms in extended models, syntactic definitions of symbols may use e.g. terms to define operation constants and type instances to define type constructors. In particular, equational definitions, well-founded recursive definitions of functions whose result types have unique choice [70], and class declarations are conservative. It will be seen below that, moreover, general recursive definitions over types of a class of domains and subtype definitions are conservative, and inductive datatype definitions as well as primitive recursive function definitions on such datatypes are conservative over base theories already containing the natural numbers.

4 Subtyping

For convenience in both writing and reading specifications, HASCASL, like CASL, features coercive subtyping. That is, basic types may be declared to be subtypes of (possibly composite) types; e.g.

types $Nat < Int$;
$Inj < Int \to Int$

declares *Nat* to be a subtype of *Int*, and *Inj* a subtype of $Int \to Int$ (say, of injective functions). The mutual subtype relation, i.e. type isomorphism, is expressed by ‘=’. Semantically, subtype relations are realised by coercion functions which are omitted in the notation. Thus, terms of the subtype may be used in terms whenever terms of the supertype are expected. Coercion functions for directly declared subtype relations are required to be injective;

however, coercion functions for inferred subtype relations as discussed further below may fail to be injective. Coercion functions are required to be coherent and compatible with overloading; this will be made more precise below. For $s < t$, one has a partial downcast operation $_\ as\ s : t \to ?s$, defined precisely on the image of s in t, and an elementhood predicate $_ \in s$ on t, defined as the predicate $\lambda\, x : t \bullet def\ x\ as\ s$. Upcasting of terms may be achieved by explicit type annotations of terms in the form $\alpha : t$.

Subtype relations may also be given polymorphically, i.e. basic type constructors may be declared to be subtypes of pseudotypes. E.g.

var	a: *Type*
type	*NonrepList* a < *List* a

declares a type constructor *NonrepList* (say, of nonrepetitive lists) such that instances of *NonrepList* are subtypes of instances of *List*. This is briefly expressed by saying that *NonrepList* is a subtype of *List* (which may be declared in the form *NonrepList* < *List*). For the built-in type constructors, we have the subtype relation

$$_ \to _ \quad < \quad _ \to ?_\,.$$

A total λ-abstraction is equal to the downcast of the corresponding partial λ-abstraction to the total function type.

A type constructor F may be declared to be *covariant* or *contravariant*, where the former means that $s < t$ implies $F\,s < F\,t$ and the latter that $t < s$ implies $F\,s < F\,t$. The absence of co- or contravariance is called *non-variance*, while the combination of contravariance and covariance is referred to as *invariance*. Covariance or contravariance of a type constructor are indicated by adding the *variance annotation* $+$ or $-$, respectively, to the corresponding constructor kind or to type variable declarations (similar ideas appear already in [10]; the notation used here is the one applied also e.g. in [1]). E.g., the list constructor might sensibly be declared to be covariant by writing

var	a: +*Type*
type	*List* a: *Type*

(or shortly *List*: +*Type* $\to$ *Type*). We do not provide dedicated syntax for invariance; however, invariant type constructors may arise through redeclaration, unions of specifications, or instantiations of parametrised specifications.

A typical example of a contravariant type constructor argument is the function type constructor: if a is a subtype of b, then $b \to c$ is, via function restriction,

a subtype of $a \to c$. Explicitly, the built-in type constructors have kinds

$$_\times_ :\ +Type \to +Type \to Type,$$
$$_\to?_,\ _\to_ :\ -Type \to +Type \to Type.$$

User-declared contravariant type constructors will mostly be related to function types in some way. E.g. one might choose to generalise the type of injective functions to a type constructor, declared by

vars $a : -Type;\ b : +Type$
type $Inj\ a\ b < a \to b$

(this must indeed be declared rather than inferred; e.g. the above variance declaration would not be sensible for a type of surjective functions).

It is possible to impose *subtyping constraints* on type variables in the form $a < t$, where a is a type variable and t is a pseudotype (similar features are present in the programming language O'Haskell [55]). For instance, the most general way to declare the *twice* function is

vars $a : Type;\ b < a$
op $twice : (a \to? b) \to (a \to? b)$

The effect is that the polymorphic profile of *twice* is annotated with the subtyping constraint $b < a$. Similarly, subtyping constraints may be imposed on polymorphic axioms. It is *not* presently allowed to impose subtyping constraints on type variables appearing as arguments in the declaration of type constructors (since this would ultimately require the introduction of 'dependent kinds'); e.g. in the above context, the declaration **type** $F\ a\ b$ would be illegal. Instances of polymorphic operations with subtyping constraints may be formed only for types that satisfy the constraints (satisfaction of constraints is decidable by means of the syntax-directed set of subtyping rules given in Appendix D).

Polymorphic operation constants introduced by the keyword **op** are required to be *coherent* under subtyping, i.e. the polymorphic instance for a subtype is required to be mapped to the instance for the supertype under the coercion function, while operators introduced by means of **fun** or **pred** are regarded as non-coherent.

Remark 33 Polymorphic predicates and functions that do not look into the structure of their type arguments, in particular typical polymorphic programs, will be coherent, while polymorphic functions or predicates involving e.g. quantification or equality will often fail to be so. Care should be taken not to accidentally declare functions of the latter kind by **op**, since this will lead to

inconsistent specifications.

Subtypes may be *defined* by means of a predicate on the supertype; e.g.

vars $a, b\colon Type$
type $Inj\ a\ b = \{f : a \to b \bullet \forall x, y\colon a \bullet f(x) = f(y) \Rightarrow x = y\}$

The total function type has the obvious subtype definition built in.

Formally, the subtype relation is a relation between type constructors and pseudotypes. In particular, it is *not* possible to declare composite types to be subtypes of others, nor to declare a subtype relation only for certain instances of a type constructor, e.g. by declaring $NonrepList\ a < List\ a$ only for $a\colon Ord$. A type constructor may be declared a subtype only of pseudotypes of the same raw kind (see below); if a type constructor F is introduced by means of a subtype declaration $F < t$ (such as *NonrepList* above) and no other kind is declared for F within the same basic specification, then F implicitly inherits the kind of t.

In the meta-theory, we denote the subtype relation by $\le$. This relation is extended to two preorders $\le$ and $\le_*$ on pseudotypes, respectively representing injective and general coercion as suggested in [28], by rules given further below. A typical case where coercions are in general non-injective is coercion by function restriction in subtype relations $b \to c \le a \to c$ for $a \le b$. Consequently, application of contravariant type constructors in general is assumed to weaken $\le$ to $\le_*$. Signatures are implicitly *embedding closed* [77], i.e. the profiles associated to a given operation constant name are upclosed under $\le_*$ in the sense that $f : s$ in Σ and $s \le_* t$ implies $f : t$ in Σ.

The set $\{\pm, +, -, \cdot\}$ of variance annotations, where $\pm$ indicates invariance and $\cdot$ is a placeholder denoting non-variance, will henceforth be denoted by $\mathcal{V}$. The set $\mathcal{V}$ is ordered by taking $\pm$ and $\cdot$ to be the smallest and the greatest element, respectively, and $+$ and $-$ to be mutually incomparable.

The subkind relation is extended by a *variance rule*

$$\frac{}{\mu Kd_1 \to Kd_2 \le_K \nu Kd_1 \to Kd_2} \quad (\mu, \nu \in \mathcal{V}, \mu \le \nu),$$

as well as analogous versions of the subkinding rule for constructor kinds (Fig. 2) for covariant and contravariant constructor kinds. The full set of subkinding rules can be found in Appendix A.

Unlike class restrictions, variance annotations are retained in *raw kinds*. This affects the admissibility of kind assignments for type constructors (Sect. 3.1); however, we syntactically relax the previous restrictions as follows: A redecla-

ration of a type constructor F may omit variances present in the raw kind of the previous declarations and also introduce new variances; these variances are then implicitly combined, and all kinds of F are modified to match the arising raw kind. E.g. the declaration of product types as generic instances of the class *Ord* in Sect. 3.1 is indeed legal, and declares the kind $+Ord \to +Ord \to Ord$ for $_ \times _$ (since $_ \times _$ has the built-in kind $+Type \to +Type \to Type$). Similarly, declaring a type to be of kind $+Cl_1 \to Cl_2 \to Cl_3$ and also of kind $Cl_1 \to -Cl_2 \to Cl_3$ results in the kind $+Cl_1 \to -Cl_2 \to Cl_3$. In the same way, different kinding declarations are reconciled in signature unions. No attempt is made to resolve conflicting variance annotations in left nested occurrences of $\to$; e.g. it is not possible to combine kindings $t\colon (+C_1 \to C_2) \to C_3$ and $t\colon (D_1 \to D_2) \to D_3$ for classes C_i, D_i. Similar relaxations apply to raw kinds of classes. Formally, we call two kinds Kd_1, Kd_2 *top-level compatible* if $Kd_i = \mu^i_1 Kd'_1 \to \cdots \to \mu^i_n Kd'_n \to Kd'$ for suitable kinds Kd'_i, Kd' and $\mu^i_j \in \mathcal{V}$, $i = 1, 2$. We admit redeclarations of types and new subclass declarations $Cl < Kd$ for existing classes Cl, provided that the associated raw kinds are top-level compatible with the previous raw kinds; the new raw kind is then the infimum of the newly declared and the previous raw kinds under the subkind relation.

The kinding rules for pseudotypes now require type contexts allowing variance-annotated type variables written in the form $a\colon +Kd$ or $a\colon -Kd$, respectively. Such variance annotations are called *outer variances*. Outer variances may appear also in type variable declarations in pseudotypes and in the actual HasCasl syntax, as already illustrated in the examples above, the effect being a variance declaration for type constructors and type synonyms declared using these type variables. Variance declarations for type synonyms are well-formed only if the associated pseudotypes are kindable by the extended kinding rules below.

The extended kinding rules concern kinding judgements $\Theta \rhd t : Kd$, with Θ a context of variance-annotated type variables, which mean that the pseudotype t depends on the variables in Θ with the indicated variance. (Strictly speaking, pseudotype formation depends also on the declared subtype constraints, but only in the sense that $\lambda\, a : Kd \bullet t$ *is ill-formed if a appears in a subtype constraint*, see above.) In type formation, only covariant or non-variant type variables can be introduced (so that e.g. the pseudotype $\lambda\, a\colon -Type \bullet a$ is ruled out). The application rule for type constructors is split into three rules

$$\frac{\Theta \rhd t : Kd_1 \qquad \Theta \rhd s : +Kd_1 \to Kd_2}{\Theta \rhd s\; t : Kd_2} \qquad \frac{\Theta^{-1} \rhd t : Kd_1 \qquad \Theta \rhd s : -Kd_1 \to Kd_2}{\Theta \rhd s\; t : Kd_2} \qquad \frac{\Theta^{0} \rhd t : Kd_1 \qquad \Theta \rhd s : Kd_1 \to Kd_2}{\Theta \rhd s\; t : Kd_2}$$

where the contexts Θ^{-1} and Θ^0 denote Θ with all outer variances reversed or

removed, respectively. For type abstraction, one has rules

$$\frac{\begin{array}{c}\Theta, a : \mu Kd_1 \rhd t : Kd_2 \\ Kd_3 \leq_K Kd_1\end{array}}{\Theta \rhd \lambda\, a : \mu Kd \bullet t : \nu Kd_3 \to Kd_2}\ (\mu \leq \nu \text{ in } \mathcal{V}).$$

By the above rules, e.g. the pseudotype $\lambda\, a \colon +\mathit{Type} \bullet \mathit{Pred}\ (\mathit{Pred}\ a)$ is of kind $+\mathit{Type} \to \mathit{Type}$, while the pseudotype $\lambda\, a \colon +\mathit{Type} \bullet \mathit{Pred}\ a$ fails to be well-formed. The full set of kinding rules is recorded in Appendix B. By induction over derivations, one shows that for t, Θ, the set $\{Kd \mid \Theta \rhd t : Kd\}$ is upwards closed w.r.t. subkinding.

Remark 34 It is not in general the case that pseudotypes have smallest kinds w.r.t. the subkind relation. E.g. the user might sensibly assign the additional kind $+\mathit{Ord} \to +\mathit{Ord} \to \mathit{Ord}$ to the product type constructor $\times$; any lower bound of that kind and the built-in kind $+\mathit{Type} \to +\mathit{Type} \to \mathit{Type}$ of $\times$ would then be a subkind of $+\mathit{Type} \to +\mathit{Type} \to \mathit{Ord}$ and hence cannot be expected to be a kind for $\times$. However, the subkinding rule for variances given above introduces a non-trivial ordering also on raw kinds, and the following proposition shows that every pseudotype has a *smallest raw kind*.

Lemma 35 *For all kinds Kd_1 and Kd_2, $Kd_1 \leq_K Kd_2$ implies $\mathsf{raw}(Kd_1) \leq_K \mathsf{raw}(Kd_2)$.*

PROOF. Induction over the derivation of $Kd_1 \leq_K Kd_2$. □

Proposition and Definition 36 *For every pseudotype t in type context Θ, the set $\{\mathsf{raw}(Kd) \mid \Theta \rhd t : Kd\}$ has a smallest element, called the* raw kind *of t in type context Θ.*

(The raw kind may be calculated by recursion along the structure of t.)

PROOF. Induction over the structure of t. The cases for type constructor and variable introduction, as well as type abstraction, are straightforward by Lem. 35. In the cases for type application, the fact is needed that $\nu Kd_2 \to Kd_3 \leq_K \mu Kd_1 \to Kd_4$ implies $Kd_3 \leq_K Kd_4$, which follows from the syntax-directed set of subkinding rules of Appendix C. □

The subtyping relations $\leq$ and $\leq_*$, ranged over by the metavariable $\sqsubseteq$, are defined by the rules of Fig. 7 (a syntax-directed version of the system is given in Appendix D). Subtyping judgements $\Theta; \Lambda \rhd s \sqsubseteq t$ in type context Θ depend on a context Λ of declared subtype constraints of the form $a \leq t$, with a a type variable and t a type in context Θ. Here, Θ is a simple type context without

outer variances. The formal difference between the two subtype relations lies in the contravariant application rule, which applies only to $\leq_*$. The subtyping rules assume that all occurring types are well-formed, i.e. kindable in the given context (in particular, the rule for abstractions assumes that a is not mentioned in Λ). The phrase '$F \leq t$ in Σ' means that the type constructor F is declared to be a subtype of t in the signature. Application of the built-in type constructors $\rightarrow ?$ etc. is covered by the application rules for arbitrary pseudotypes.

$$\frac{a \leq t \text{ in } \Lambda}{\Theta; \Lambda \rhd a \leq t} \quad \frac{F \leq t \text{ in } \Sigma}{\Theta; \Lambda \rhd F \leq t} \quad \frac{\Theta; \Lambda \rhd s \leq t}{\Theta; \Lambda \rhd s \leq_* t} \quad \frac{}{\Theta; \Lambda \rhd t \leq t} \quad \frac{\Theta; \Lambda \rhd s \sqsubseteq t \quad \Theta; \Lambda \rhd t \sqsubseteq u}{\Theta; \Lambda \rhd s \sqsubseteq u}$$

$$\frac{\Theta \rhd t : +Kd_1 \rightarrow Kd_2 \quad \Theta; \Lambda \rhd s_1 \sqsubseteq s_2}{\Theta; \Lambda \rhd t\ s_1 \sqsubseteq t\ s_2} \quad \frac{\Theta \rhd t : -Kd_1 \rightarrow Kd_2 \quad \Theta \rhd s_2 \leq_* s_1}{\Theta \rhd t\ s_1 \leq_* t\ s_2} \quad \frac{\Theta; \Lambda \rhd t_1 \sqsubseteq t_2}{\Theta; \Lambda \rhd t_1\ s \sqsubseteq t_2\ s}$$

$$\frac{\Theta, a : Kd; \Lambda \rhd t \sqsubseteq s}{\Theta; \Lambda \rhd \lambda\, a : \mu Kd \bullet t \sqsubseteq \lambda\, a : \mu Kd \bullet s}\ (\mu \in \mathcal{V})$$

Fig. 7. Subtyping rules for pseudotypes (with $\sqsubseteq \in \{\leq, \leq_*\}$)

Lemma 37 *For pseudotypes t_1 and t_2, $\Theta; \Lambda \rhd t_1 \leq_* t_2$ implies that t_1 and t_2 have the same raw kinds.*

Semantically, the interpretation of subtyping is determined by an extension of the translation of polymorphic HasCasl into basic HasCasl to signatures with subtyping, defined as follows.

Definition 38 A *polymorphic* HasCasl *signature* Σ *with subtyping* is defined by extending the notion of polymorphic HasCasl signature (Defn. 16) in the way indicated above: there is additional data in the shape of the subtyping relation $\leq$ between type constructors and pseudotypes, and a *coherence* predicate on the set of polymorphic operations (see above). Moreover, polymorphic operations and axioms are annotated with sets of subtyping constraints of the form described above. For semantic purposes, we admit also subtyping constraints of the form $a \leq_* t$ (such constraints are never generated by user declarations). The restrictions listed above apply, in particular embedding closure. Besides the user-declared symbols, Σ implicitly contains polymorphic operations (assumed to be different from all user-declared oper-

ations)

$$up : \forall a, b : \mathit{Type}; a \leq_* b \bullet a \to b$$
$$down : \forall a, b : \mathit{Type}; a \leq b \bullet b \to ?a.$$

Similarly, *morphisms* of such signatures are defined by extending the definition of morphism of polymorphic signatures. Signature morphisms map only the user-declared symbols (not the above implicit operations). They are required to preserve coherence and the subtype relation $\leq$, the latter in the sense that subtype declarations are mapped to derivable subtyping judgements. We impose that overloading of symbols is preserved [53], a condition which thanks to embedding closure reduces to the requirement that identically named constants $c : s$ and $c : t$ are mapped to identically named constants whenever $s < t$. Moreover, we require that raw kinds of classes and type constructors are preserved up to top-level compatibility (by Lem. 35 and Prop. 36, it follows already from preservation of the kinding and subclass relations that raw kinds can only decrease w.r.t. the subkind relation).

Over a polymorphic HasCasl signature Σ with subtyping, we define two kinds of sentences: *explicit coercion* sentences are just the expected polymorphic sentences over Σ, including the built-in symbols, with instances of polymorphic operations admitted only for types satisfying the associated subtyping constraints. *Implicit coercion* sentences additionally may use the above-mentioned subtyping mechanisms, i.e. terms of a subtype can appear wherever terms of a supertype are expected, and downcasts $_\ \mathit{as}\ s$ and elementhood $_ \in s$ may be used; however, implicit coercion sentences cannot use the built-in symbols *up* and *down*. Implicit coercion sentences are used in actual specifications, while explicit coercion sentences serve only semantic purposes. Implicit coercion sentences are translated into explicit coercion sentences by

- inserting *up* where terms of a subtype are used in places where terms of the supertype are expected;
- replacing uses of $\in$ with its definition in terms of *down* and definedness;
- replacing downcasts $_\ \mathit{as}\ s$ with applications of *down*.

The translation of polymorphic signatures into basic signatures is extended by associating to a polymorphic signature Σ with subtyping a basic *theory* $\mathsf{B}(\Sigma)$ as follows. The signature of $\mathsf{B}(\Sigma)$ is defined as before (Defn. 18), except that subtyping constraints are taken into account: a polymorphic operation constant $f : \forall a_1 : Kd_1, \dots, a_n : Kd_n; \Lambda \bullet t$ is instantiated only to those type instances $s_1, \dots, s_n$ that satisfy every subtyping constraint $a_i \sqsubseteq t$ in Λ (with $\sqsubseteq \in \{\leq, \leq_*\}$) in the sense that $(); () \rhd s_i \sqsubseteq t[s_1/a_1, \dots, s_n/a_n]$ is derivable in Σ, using for generalised types the additional rule that $(u.\,\phi) \leq (u.\,\psi)$ if $\forall x : u \bullet \phi\ x \Rightarrow \psi\, x$ is derivable. The axioms of $\mathsf{B}(\Sigma)$ are obtained by translating, in the way described below, the following explicit coercion sentences over Σ:

- coercion from s to t is injective if $s \leq t$, with *down* as a partial left inverse:

$$\forall a, b : \mathit{Type}; a \leq b \bullet\ (\forall x : a \bullet \mathit{down}\ ((\mathit{up}\ x) : b) = x) \wedge \\ \forall y : b \bullet \mathit{def}\ (\mathit{down}\ y) : a\ \Rightarrow \mathit{up}\ ((\mathit{down}\ y) : a) = y$$

- subtyping is *coherent*, i.e. coercion functions compose and coercion from a type into itself is the identity:

$$\forall a : \mathit{Type} \bullet \forall x : a \bullet ((\mathit{up}\ x) : a) = x \text{ and} \\ \forall a, b, c : \mathit{Type}; b \leq_* c, a \leq_* b \bullet \forall x : a \bullet \mathit{up}\ ((\mathit{up}\ x) : b) = (\mathit{up}\ x) : c.$$

- overloading of operations is compatible with coercion, i.e. for each type context $\Theta = (a_1 : \mathit{Kd}_1; \ldots; a_n : \mathit{Kd}_n)$, each polymorphic operation $c : \forall\Theta; \Lambda \bullet s$, and each type t such that $\Theta; \Lambda \rhd s \leq_* t$, there is an axiom

$$\forall\Theta; \Lambda \bullet \mathit{up}\ (c : s) = c : t$$

(where the profile $c : \forall\Theta; \Lambda \bullet t$ is in Σ by embedding closure);
- correspondingly flagged polymorphic operations are coherent w.r.t. subtyping: if $f : \forall b_1 : \mathit{Kd'}_1, \ldots, b_m : \mathit{Kd'}_m \bullet s$ is a coherent polymorphic operation, and for $i = 1, \ldots, m$, t_i and u_i are types of kind $\mathit{Kd'}_i$ such that $\Theta; \Lambda \rhd s[t_1/b_1, \ldots, t_m/b_m] \leq_* s[u_1/b_1, \ldots, u_m/b_m]$, then there is an axiom

$$\forall\Theta; \Lambda \bullet f[u_1, \ldots, u_m] = \mathit{up}\ f[t_1, \ldots, t_n];$$

- the built-in subtype relations have the expected coercion functions; i.e.

$$\forall a, b, c, d : \mathit{Type}; c \leq_* a; d \leq_* b \bullet \\ (\forall f : a \to b \bullet (\mathit{up}\ f) : (a \to ?b) = \lambda x : a \bullet\ f\, x) \wedge \\ (\forall f : a \to ?d \bullet (\mathit{up}\ f) : (c \to ?b) = \lambda x : c \bullet\ \mathit{up}\ (f\ (\mathit{up}\ x))) \wedge \\ \forall x : c; y : d \bullet (\mathit{up}\ (x, y)) : a \times b = (\mathit{up}\ x, \mathit{up}\ y)$$

(note that η does not apply to the right hand side in the first equation, since f has the wrong type).

Finally, explicit coercion sentences are translated into collections of sentences over $\mathsf{B}(\Sigma)$ in the same way as in Defn. 18, with instances restricted to those satisfying the given subtyping constraints. A model of Σ is a model of $\mathsf{B}(\Sigma)$, and such a model satisfies a Σ-sentence if it satisfies all its instances.

Remark 39 The subtyping axioms above imply that the subtype of total functions contains all total functions that live in the partial function type (see Defn. 9) and that co-contravariant subtype relations for total function types have the right coercion functions, i.e.

$$\forall f : a \to d \bullet (\mathit{up}\ f) : (c \to b) = \lambda\, x : c \bullet !\, \mathit{up}\ (f\ (\mathit{up}\ x)).$$

Remark 40 The presence of the *down* operation implies that subtypes $a \leq b$ are regular subobjects in the categorical models [71].

5 Datatypes

HasCasl supports recursive datatypes in the same style as in Casl. To begin, an unconstrained family of datatypes t_i is declared along with its constructors $c_{ij} : t_{ij1} \to \ldots \to t_{ijk_{ij}} \to t_i$ by means of the keyword **type** in the form

$$\begin{array}{ll} \textbf{type} & t_1 ::= \; c_{11}\; t_{111}\; \ldots\; t_{11k_{11}} \;\mid\; \ldots \;\mid\; c_{1m_1}\; t_{1m_11}\; \ldots\; t_{1m_1k_{1m_1}} \\ & \ldots \\ & t_n ::= \; c_{n1}\; t_{n11}\; \ldots\; t_{n1k_{n1}} \;\mid\; \ldots \;\mid\; c_{nm_n}\; t_{nm_n1}\; \ldots\; t_{nm_nk_{nm_n}} \end{array}$$

Here, the t_i may be patterns of the form $C\ a_1\ \ldots\ a_r$, where C is a type constructor and the a_i are type variables, so that C is declared as a polymorphic type. The t_{ijl} are types in the context determined by the C and the a_p. Optionally, selectors $s_{ijl} : t_i \to ?t_{ijl}$ may be declared by writing $(s_{ijl} :?t_{ijl})$ in place of t_{ijl}. All this is just syntactic sugar for the corresponding declarations of types and constants, and equations stating that selectors are left inverse to constructors.

Data types may be qualified by a preceding **free** or **generated**. The **generated** constraint introduces an induction axiom; this corresponds roughly to term generatedness ('no junk'). The **free** constraint ('no junk, no confusion') instead introduces an implicit fold operation, which implies both induction and a primitive recursion principle. If one of these constraints is used, then recursive occurrences (in the t_{ijl}) of a type constructor C being declared are restricted to the pattern $C\ a_1\ \ldots\ a_r$ appearing on the left hand side; i.e. HasCasl does not support polymorphic recursion. If a **free** constraint is used, then additionally recursive occurrences of the types being declared are required to be strictly positive w.r.t. function arrows, i.e. occurrences in the argument type of a function type are forbidden.

In more detail, the semantics of the constraints is as follows.

5.1 Generated types

For types t_i as above that have only types t_j from the same declaration and types from the local environment as arguments of constructors, the induction axiom states that for any family of predicates $P_i : \mathit{Pred}\ t_i$, called the *induction predicates*, the premise of the induction principle implies that $\forall x : t_i.\, P_i(x)$ for

all i. Here, the premise of the induction principle expresses that the induction predicates are closed under the constructors in the usual sense. Note that the induction axiom is a higher-order reformulation of the corresponding sort generation constraint in CASL. Unlike in CASL, the induction axioms are however expressible in HASCASL, i.e. generation constraints in HASCASL are just syntactic sugar.

For constructors with composite argument types, the notion of closedness of predicates under the constructor requires extending the induction predicates to *extended induction predicates* P_s on composite types s, as follows.

- *Partial function spaces*:

$$P_{s \to ?t} f \iff \forall x : s.(P_s x \wedge \operatorname{def} f(x)) \Rightarrow P_t f(x).$$

- *Total function spaces*: $P_{s \to t}$ is the restriction of $P_{s \to ?t}$ to $s \to t$.
- *Product types*:

$$P_{s \times t}(x, y) \iff P_s x \wedge P_t y.$$

- For types s from the *local environment*, P_s is taken to be constantly true.
- *Applications* $D\ s_1\ \ldots\ s_q$ of a type constructor D from the local environment to types $s_1, \ldots, s_q$, where at least one s_j contains a recursive occurrence of the t_j: extended induction predicates for such types are required to be closed under all operations with result type $D\ s_1\ \ldots\ s_q$ (which are necessarily newly arising instances of polymorphic operations). Note that in general, extended induction predicates are not uniquely defined by this requirement. Examples follow below.

Remark 41 If a type constructor D from the local environment has a generation constraint, then of course the closedness requirement on extended induction predicates for applications of D is equivalent to closedness under the operations in the constraint. However, the induction axiom also makes sense if D has no generation constraint; it then states essentially that the types being declared are generated from the reachable part of D. Note that extended induction predicates do not appear in the conclusion of the induction axiom, so that the latter does not imply a sort generation constraint for D.

Generally, every HASCASL specification, in particular every datatype declaration, has a term model [71], and the induction axiom induced by a generatedness constraint is satisfied in the term model. However, we stress that the induction axiom does *not* imply that elements of a generated datatype whose constructors have functional arguments are reachable by the constructors and λ-abstraction. In particular, concerning inhabitants of functional types, the induction axioms only require that these map reachable elements to reachable elements — they need not be themselves reachable. Specifically, a standard interpretation of functional types (i.e. using the full function space, which cannot be term generated for infinite types) is not precluded.

Finally, note that, due to the flexibility of interpretation of higher types in Henkin models, the higher-order reformulation of generation constraints in HasCasl is weaker than the corresponding generation constraint in Casl, and in particular does not exclude non-standard models. However, proof-theoretically, this difference disappears — at least if the standard Casl proof system with the usual finitary induction rule is used. Only if stronger (e.g. infinitary) forms of induction are used, the difference becomes relevant. It also becomes relevant for monomorphicity. A specification is called *monomorphic* if all its models are isomorphic. Due to possible non-standard interpretations of higher types, even free datatypes (see below) are not monomorphic in HasCasl, although they are monomorphic in Casl.

Example 42 The following datatype declaration (to be extended by a precise specification of equality on the declared types) might form part of a specification of finite systems with unordered branching:

generated type $\mathit{Container}\ a ::= \mathit{empty} \mid \mathit{add}\ a\ (\mathit{Container}\ a)$
generated type $\mathit{Sys}\ b ::= \mathit{node}\ b\ (\mathit{Container}\ (\mathit{Sys}\ b))$

The induction axiom for *Container* is as in Casl; the induction axiom for *Sys b* is as follows.

$$\left.\begin{array}{l}(\forall x : b; s : \mathit{Container}\ (\mathit{Sys}\ b) \bullet \\ \qquad\qquad Q\ s \Rightarrow P\ (\mathit{node}\ x\ s)) \wedge \\ Q\ \mathit{empty} \wedge \\ (\forall s : \mathit{Container}\ (\mathit{Sys}\ b); t : \mathit{Sys}\ b \bullet \\ \qquad\qquad (Q\ s \wedge P\ t) \Rightarrow Q\ (\mathit{add}\ t\ s))\end{array}\right\} \Rightarrow \forall t : \mathit{Sys}\ b \bullet P\ t.$$

As an example with functional constructors, consider a datatype of at most countably branching trees,

generated type $\mathit{CTree}\ a ::= \mathit{leaf}\ a \mid \mathit{branch}\ (\mathit{Nat} \rightarrow ?\ \mathit{CTree}\ a)$

(with the type *Nat* of natural numbers declared elsewhere), which for *CTree* gives rise to the induction axiom

$$\left.\begin{array}{l}(\forall x : a \bullet P\ (\mathit{leaf}\ x)) \wedge \\ (\forall f : \mathit{Nat} \rightarrow ?\mathit{CTree}\ a \bullet \\ \quad (\forall n : \mathit{Nat} \bullet \mathrm{def}\ f\ n \Rightarrow P\ (f\ n)) \Rightarrow P\ (\mathit{branch}\ f))\end{array}\right\} \Rightarrow \forall t : \mathit{CTree}\ a \bullet P\ t.$$

The semantics of free datatypes is determined by a fold operation, i.e. free datatypes are explicitly axiomatised as initial algebras. As indicated above, negative occurrences of the types being declared are forbidden in declarations of free types, i.e.

free type $L ::= \mathit{abs}\ (L \to L)$

is illegal, while

free type $\mathit{Tree}\ a\ b ::= \mathit{leaf}\ b \mid \mathit{branch}\ (a \to \mathit{Tree}\ a\ b)$

is allowed. Thanks to this restriction, we can set about interpreting free datatypes as initial algebras for functors.

To begin, a declaration of datatypes $t_1, \ldots, t_n$ as above can be regarded as a fixed point declaration for a family $F = (F_1, \ldots, F_n)$ of n-argument type constructors; here, alternatives $A \mid B$ are replaced by sums $A + B$, using a built-in declaration of sum types as in Sect. 3.1. The constructors of the t_i can then be gathered into *structure maps* $c_i : F_i\ t_1\ \ldots\ t_n \to t_i$.

We then extend F to a functor, where we view a functor as mapping types to types and functions to functions as in Fig. 4. The action of F on maps is defined by recursion over the structure of F, with the standard clauses for sums, products, function types (where only positive positions appear), and constant types, i.e. types from the local environment. The remaining case in the recursion are types $D\ s_1\ \ldots\ s_q$, where D is a type constructor from the local environment. Here, we have to require that the functorial action of D is determined by its specification; that is, the free type is well-formed only if D belongs to the class $\mathit{Functor}_q$, a built-in specification of functors with q arguments that generalises the specification $\mathit{Functor} = \mathit{Functor}_1$ of Fig. 4. (For practical purposes, q can be restricted to small values.)

If the t_i are parametrised over type variables $a_i : \mathit{Type}$, then F is parametrised in the same way. If the a_i appear only in functorial positions, so that the t_i depend functorially on the a_i, then corresponding instances of $\mathit{Functor}_q$ are derived automatically. (Note that this means that the average user will never actually see the classes $\mathit{Functor}_q$ in practice, as instances are generated and exploited automatically for typical sequences of nested datatype declarations.) In order to keep the language design manageable, functorial dependence on variables of higher kinds, although technically possible, is not supported. See [72] for details on the functor mechanism.

The fold operations $fold_i$ for the t_i then have the polymorphic types

$$\begin{aligned} fold_i : \forall b_1, \dots, b_n : Type \bullet (F_1\ b_1\ \dots\ b_n \to b_1) \to \dots \\ \to (F_n\ b_1\ \dots\ b_n \to b_n) \to t_i \to b_i. \end{aligned}$$

(In practice, if F_i is a sum type arising from alternatives, then the argument of type $F_i\ b_1\ \dots\ b_n$ is decomposed into several functions, one for each component of the sum.) The defining property of the fold operations states that, for $b_1, \dots, b_n : Type$, $f_i : t_i \to b_i$, and $d_i : F_i\ b_1\ \dots\ b_n \to b_i$, $i = 1, \dots, n$,

$$f_i = fold_i\ d_1\ \dots\ d_n \text{ for all } i \quad \text{iff} \quad d_i \circ (F_i\ f_1\ \dots\ f_n) = f_i \circ c_i \quad \text{for all } i;$$

$$\begin{array}{ccc} F_i\ t_1\ \dots\ t_n & \xrightarrow{F_i\ f_1\ \dots\ f_n} & F_i\ b_1\ \dots\ b_n \\ \downarrow c_i & & \downarrow d_i \\ t_i & \xrightarrow{f_i} & b_i \end{array}$$

i.e. the $fold_i\ d_1\ \dots\ d_n$ constitute the unique F-algebra morphism from the t_i into the F-algebra given by the d_i. Thus, the c_i are determined as the initial F-algebra.

Remark 43 The above requires two warnings. To begin, although we define datatypes as internal initial algebras, they are not in general monomorphic; e.g., the standard definition of the naturals as a free datatype admits non-standard models. This is due to the Henkin semantics — the set of functions to which the fold operation applies does not have a fixed interpretation.

Secondly, unlike in first order CASL, the meaning of **free type** does not coincide with that of the corresponding structured free extension **free { type** ... **}**. The difference is that a free extension also requires all newly arising function types to be freely term generated, which has the undesirable effect of precluding any further function definitions for these types (going beyond existing λ-terms).

Example 44 Consider the following free datatype definitions.

vars a, b: *Type*
free type $List\ a ::= nil \mid cons\ a\ (List\ a)$
free type $Tree\ a\ b ::= leaf\ a \mid branch\ (b \to List\ (Tree\ a\ b))$

For the type constructor *List*, an instance *List* : *Functor* is derived, with *map* defined in the standard way. For *Tree*, we obtain an operation

$$fold : (a \to c) \to ((b \to List\ c) \to c) \to Tree\ a\ b \to c,$$

CTree polymorphic over $c : Type$. This operation is axiomatised as being uniquely determined by the equations

$$\begin{aligned} fold\ f\ g\ (leaf\ x) &= f\ x \\ fold\ f\ g\ (branch\ s) &= g\ (map\ (fold\ f\ g) \circ s). \end{aligned}$$

Remark 45 From the fold operation, one obtains also a primitive recursion principle in the standard way (i.e. by means of a simultaneous recursive definition of the identity). From the latter, in turn, we obtain as a special case a case operator, denoted in the form

$$case\ x\ of\ c\ y_1\ \dots\ y_l \to f\ y_1\ \dots\ y_l \mid \dots$$

Moreover, free types are generated, i.e. satisfy the induction axiom of Sect. 5.1. In the following, we regard sum types (Fig. 3), which we denote by + in the interest of readability, as free datatypes, and in particular use the case notation for them as well.

Remark 46 Unlike in Casl, declarations of free datatypes in HasCasl are not necessarily a conservative extension of the local environment. Already the naturals may be a non-conservative extension: as discussed in Sect. 3.2, conservative extensions at the second level of the semantics essentially can only introduce names for entities already in the present signature. However, if the naturals, as well as sum types and an initial object as specified in Fig. 3, are already present, then one can show *finitely branching* free datatypes to be conservative extensions. This is done by constructing them in a similar way as in standard HOL [57,5] as inductively generated subtypes of a suitable universal type, with some modifications required due to the fact that HasCasl does not impose unique choice (Sect. 2.2) — essentially, the universal type is a type of trees, represented as partial maps from paths to values (rather than as sets of (node, value)-pairs as in HOL), and annotated explicitly with finite sizes in order to inherit primitive recursion from the naturals. Details are laid out in [72].

6 Recursion

Unlike Casl, HasCasl has a notion of executable specification that includes general recursion and hence possible non-termination, in the style of a strict functional programming language (as laid out in Sect. 2.3, non-strict functions can be modelled as well). This is achieved by explicitly bootstrapping a domain semantics in the style of HOLCF [64]. On the technical side, this requires some adjustments to standard domain theory in order to cope with the austerity of the internal logic; these issues are dealt with in Sect. 6.1. We then go on to

discuss initial datatypes in the arising category of domain types, and finally describe how these features are reflected by an appropriate sugaring of the HasCasl syntax.

6.1 Domain Theory in the Internal Logic

We now recast the basics of standard domain theory, phrased in terms of chain-complete partial orders, in the internal logic of HasCasl. The main difficulty here is not so much the intuitionistic aspect (the study of domain theory in toposes goes back at least to [65]), but the fact that due the the absence of unique choice (Sect. 2.2), we can no longer e.g. define the value at x of the supremum of a chain of partial functions f_i as 'the value (if any) eventually assumed by the $f_i(x)$'. Rather, we have to require existence of suprema of chains in the lifting $?a$ of a domain a; for this purpose, we assume given in this section a type *Nat* of natural numbers.

Definition 47 A partial order a with ordering $\sqsubseteq$ is called a *complete partial order (cpo)* if the lifted type $?a$, equipped with the ordering

$$x \sqsubseteq y \iff (\mathit{def}\, x() \Rightarrow x() \sqsubseteq y()),$$

has suprema of chains and a bottom element. We call chains in $?a$ *partial chains*, as opposed to *total chains*, i.e. chains in a in the usual sense. We say that a cpo a is *pointed* (or a *cppo*) if a has a bottom element. We say that a type a is a *flat cpo* if a is a cpo when equipped with the discrete ordering.

Cpo's can be specified as a class in HasCasl as shown in Fig. 9; the specification imports, besides the natural numbers, a specification of partial orders (Fig. 8), containing in particular the definition of the induced ordering on lifted types. In the discussion below, we denote suprema of (total or partial) chains by $\bigsqcup$.

Remark 48 Note that the notation for the ordering is changed from $\leq$ to $\sqsubseteq$ in Fig. 9, and the class *Ord* is renamed into *InfOrd* (information ordering) in order to allow the future declaration of the expected instances of the class *Ord*, e.g. the usual ordering on the flat cpo of natural numbers. This nicely illustrates the benefits of combining a class mechanism with Casl's structured specification constructs. In a framework without such constructs, such as Isabelle, it becomes necessary at this point to fully respecify a second copy of the class *Ord* (and indeed this is precisely what happens in Isabelle/HOLCF, where a class `po` of partial orders with ordering $\sqsubseteq$ is newly specified although Isabelle/HOL already includes a class `order` of partial orders with ordering $\leq$).

As in standard domain theory, cppo's in the above sense have least fixed points of continuous endofunctions f, constructed as suprema of (total) chains $(f^n \bot)$. The fixed point operator is denoted by Y. For properties P of $Y\ f$, one has the standard *fixed point induction* principle: if $P\ \bot$, $P\ x \Rightarrow P\ (f\ x)$, and P is *admissible*, i.e. closed under suprema of total chains, then $P\ (Y\ f)$. Both the definition of Y and the fixed point induction theorem are explicitly included in Fig. 9.

spec ORD =
  class *Ord* {**var** a : *Ord*
  fun $_ \leq _$: *Pred* $(a \times a)$
  var $x, y, z : a$
   • $x \leq x$
   • $(x \leq y \wedge y \leq z) \Rightarrow x \leq z$
   • $(x \leq y \wedge y \leq x) \Rightarrow x = y$
   }
  var a, b : +*Ord*
  type instance $a \times b$: *Ord*
  var $x, z : a;\ y, w : b$
   • $(x, y) \leq (z, w) \Leftrightarrow x \leq z \wedge y \leq w$
  type instance *Unit* : *Ord*
   • $() \leq ()$
  type instance $?a$: *Ord*
  var $x, y :?a$
   • $x \leq y \Leftrightarrow (\mathit{def}\ x() \Rightarrow x() \leq y())$

Fig. 8. Specification of partial orders

To begin, we note that under unique choice, i.e. for coarse types, the definition of cpo coincides with the usual one using total chains.

Proposition 49 *If a coarse type a has suprema of total chains, then a is a cpo in the sense of Fig. 9.*

PROOF. The supremum of a partial chain (x_i) in a is $\iota\, x : a \bullet \phi$, where ϕ states that there exists n such that (x_{i+n}) is a total chain with supremum x. The bottom element of $?a = \mathit{Unit} \to ?a$ is the unique function $\bot$ such that $\neg \mathit{def} \bot()$. □

Example 50 The above proposition implies in particular that coarse types become cpo's when equipped with the discrete ordering, so that one has the usual concept of flat cpo. There are natural examples of models where *all* types can be made into flat cpo's, but also equally natural examples demonstrating that this need not be the case.

```
spec  Recursion = {Ord with Ord ↦ InfOrd, __ ≤ __ ↦ __ ⊑ __} and Nat
      then
      class   Cpo < InfOrd {var a: Cpo
      fun     __ ⊑ __: Pred (a × a)
      op      undefined : ?a
              • ¬def (undefined : ?a)
      type    Chain a = {s: Nat →? a • ∀n: Nat • def s n ⇒ s n ⊑ s (n + 1)}
      fun     sup : Chain a →?a
      var     x :?a; s: Chain a
              • sup s ⊑ [?a] x ⇔ ∀n: Nat • s n ⊑ [?a] x
              }
      class   Cppo < Cpo {var a: Cppo; x: a
      fun     bottom : a
              • bottom ⊑ x }
      class   FlatCpo < Cpo {vars a: FlatCpo; x, y: a
              • x ⊑ y ⇒ x = y }
      vars    a, b: Cpo; c: Cppo; x, y: a; z, w: b
      type instance __ × __: +Cpo → +Cpo → Cpo
      type instance __ × __: +Cppo → +Cppo → Cppo
      type instance Unit: Cppo
      type instance Unit: FlatCpo
      type    a -c->? b = {f : a →?b • ∀s: Chain a •
                          let t = λn: Nat • f (s n) in
                          t ∈ Chain b ⇒ sup (t as Chain b) = f (sup s)}
      type    a -c-> b = {f : a -c->? b • f ∈ a → b}
      type instance __ -c->? __: −Cpo → +Cpo → Cppo
      var     f, g: a -c->? b • f ⊑ g ⇔ ∀x: a • def (f x) ⇒ f x ⊑ g x
      type instance __ -c-> __: −Cpo → +Cpo → Cpo
      var     f, g: a -c-> b • f ⊑ g ⇔ ∀x: a • f x ⊑ g x
      type instance __ -c-> __: −Cpo → +Cppo → Cppo
              • bottom[a -c-> c] = λx: a •! bottom[c]
      then    %def
      var     c: Cppo
      fun     Y : (c -c-> c) -c-> c
      var     f : c -c-> c; x: c; P: Pred c
              • f (Y f) = Y f
              • f x = x ⇒ Y f ⊑ x
              • (P bottom ∧ (∀x: c • P x ⇒ P (f x))) ⇒ P (Y f)   %implied
```

Fig. 9. Specification of the cpo structure and the fixed point operator

As a positive example, consider the quasitopos **ReRe** of reflexive relations [2], whose objects are pairs (X, R) with R a reflexive relation on the set X, and whose morphisms are relation-preserving maps. In a model over **ReRe**, the interpretation $(X_\perp, R_\perp)$ of $?a$ is obtained from the interpretation (X, R) of a by adding a new element $\perp$ to X and putting $x R_\perp \perp$ and $\perp R_\perp x$ for all $x \in X_\perp$.

It is easy to check that in this case, a indeed becomes a flat cpo, the crucial point being that the supremum operation is a relation-preserving map from the type of partial chains to $?a$.

A negative example is given by the quasitopos **PsTop** of pseudotopological spaces and continuous functions [30]. A *pseudotopological space* is given in terms of a convergence relation $\to$ between filters on a set X and points of X, subject to the requirements that $\dot{x} \to x$, where $\dot{x} = \{A \subseteq X \mid x \in A\}$, and that $\mathfrak{F} \to x$ iff for all ultrafilters $\mathfrak{U}$ finer than $\mathfrak{F}$ (i.e. $\mathfrak{F} \subseteq \mathfrak{U}$), $\mathfrak{U} \to x$. A function f between pseudotopological spaces is *continuous* if $f(\mathfrak{F}) \to f(x)$ whenever $\mathfrak{F} \to x$, where $f(\mathfrak{F}) = \{A \mid f^{-1}[A] \in \mathfrak{F}\}$. In **PsTop**, the interpretation $X_\perp$ of $?a$ is obtained from the space X interpreting a type a by adding an element $\perp$ and putting $\mathfrak{F} \to \perp$ for all filters $\mathfrak{F}$ on $X_\perp$, and $\mathfrak{F} \to x$ iff $\mathfrak{F}_X \to x$ for $x \in X$, where $\mathfrak{F}_X = \{A \cap X \mid A \in \mathfrak{F}\}$. Moreover, if types a and b are interpreted by spaces X and Y, respectively, then the function type $a \to b$ is interpreted by the space $X \to Y$, consisting of the continuous functions $f : X \to Y$, with $\mathfrak{F} \to f$ for a filter $\mathfrak{F}$ on $X \to Y$ iff, whenever $\mathfrak{G} \to x$ in X, then $ev(\mathfrak{F} \times \mathfrak{G}) \to f(x)$, where $\mathfrak{F} \times \mathfrak{G}$ denotes the filter generated by the set $\{A \times B \mid A \in \mathfrak{F}, B \in \mathfrak{G}\}$, and $ev(f, x) = f(x)$. A discrete pseudotopological space is characterised by $\mathfrak{F} \to x$ iff $\mathfrak{F} = \dot{x}$. For discrete spaces X, in particular the natural numbers object (i.e. the set $\mathbb{N}$ equipped with the discrete structure), the above definition of $\mathfrak{F} \to f$ in $X \to Y$ simplifies to $ev_x(\mathfrak{F}) \to f(x)$ for all $x \in X$, where $ev_x(f) = f(x)$.

In models over **PsTop**, no discrete space B with at least two points is a flat cpo in the above sense. The reason is that the supremum map $\sup : \mathit{Chain}\ B \to B_\perp$, where the type $\mathit{Chain}\ B$ of partial chains in B inherits its convergence relation from $\mathbb{N} \to B_\perp$ as a subspace, fails to be continuous. To see this, pick $s \in \mathit{Chain}\ B$ such that s_m is defined for some $m \in \mathbb{N}$. Let $\mathfrak{F}$ be the filter generated by the set

$$\{ev_n^{-1}[C \cup \{\perp\}] \mid C \subseteq B, n \in \mathbb{N}, s_n \in C\}.$$

One can check that $\mathfrak{F} \to s$ in $\mathit{Chain}\ B$. Thus in order for sup to be continuous, one would need $\sup(\mathfrak{F}) \to \sup s = s_m$ in $B_\perp$, whence for all $C \subseteq B$, $\sup^{-1}[C \cup \{\perp\}] \in \mathfrak{F} \iff s_m \in C$ by discreteness of B. In particular, $\sup^{-1}[\{s_m, \perp\}] \in \mathfrak{F}$, i.e. there exist $C_1, \dots, C_k \subseteq B$ and $m_1, \dots, m_k \in \mathbb{N}$ such that

$$\sup{}^{-1}[\{s_m, \perp\}] \supseteq \bigcap ev_{m_i}[C_i \cup \{\perp\}].$$

Now pick $t \in \mathit{Chain}\ B$ such that $t_{m_i} = \perp$ for all i and $\sup t \in B - \{s_m, \perp\}$; then t is contained in the right hand side of the above formula, but not in the left hand side, contradiction.

We now verify that a number of domain theoretic constructions work for our definition of cpo, as claimed by the instance declarations in Fig. 9. Partial suprema have the expected behaviour:

Lemma 51 *Let (x_i) be a partial chain in a. Then $\bigsqcup_i x_i$ is defined iff $\exists n \bullet$ def x_n.*

PROOF. The 'if' direction is trivial. Concerning 'only if', just note that

$$(\bigsqcup x_i) \ \mathit{res} \ \exists n \bullet \mathit{def} \ x_n$$

is an upper bound of (x_i). □

According to Fig. 9, a partial function between cpo's is *continuous* iff it preserves suprema of partial chains. This is equivalent to the standard definition in terms of Scott open domains of definition and preservation of suprema of total chains:

Definition 52 A predicate $P : \mathit{Pred}\ a$ is called *Scott open* if P is upclosed, i.e. $P\ x$ and $x \sqsubseteq y$ imply $P\ y$, and $P\ (\bigsqcup x_i)$ for a total chain (x_i) implies $\exists n \bullet P\ x_n$.

Proposition 53 *Let a and b be cpo's, and let $f : a \to ?b$. Then f is continuous iff the predicate $P = \lambda x : a \bullet \mathit{def}\ (f\ x)$ is Scott open, f is monotone on P, and f preserves suprema of total chains (x_i) in the sense that if $f\ (\bigsqcup_i x_i)$ is defined then there exists m such that $f\ x_m$ is defined and $f\ \bigsqcup_i x_i = \bigsqcup_i (f\ x_{i+m})$.*

PROOF. *'Only if':* If $f\ x$ is defined and $x \sqsubseteq y$, then we have a chain x_i recursively defined by $x_0 = x$ and $x_{i+1} = y$. Thus $f\ y = f\ \bigsqcup_i x_i = \bigsqcup_i f\ x_i \geq f\ x$ in $?b$, so that $f\ y$ is defined. This proves monotonicity of f and the first part of Scott openness. For the second part of the latter and preservation of suprema of total chains, let (x_i) be a total chain such that $f\ \bigsqcup_i x_i$ is defined. By continuity, $\bigsqcup_i f\ x_i$ is defined and equal to $f\ \bigsqcup_i x_i$, so that by Lem. 51, there exists m such that $f\ x_m$ is defined; then $f\ \bigsqcup_i x_i = \bigsqcup f\ x_i = \bigsqcup f\ x_{i+m}$.

'If': Let (x_i) be a partial chain. We have to prove the strong equation $f\ \bigsqcup x_i = \bigsqcup f\ x_i$. To begin, assume that $f\ \bigsqcup_i x_i$ is defined. Then $\bigsqcup_i x_i$ is defined, so that by Lem. 51, there exists m such that x_m is defined. Then (x_{i+m}) is a total chain and $\bigsqcup x_{i+m} = \bigsqcup x_i$; hence we have n such that $f\ x_{i+m+n}$ is defined and $f\ (\bigsqcup x_i) = \bigsqcup (f\ x_{i+m+n}) = \bigsqcup (f\ x_i)$.

Conversely, let $\bigsqcup (f\ x_i)$ be defined. By Lem. 51, we have n such that $f\ x_n$ is defined. Since P is upclosed, it follows that $f\ (\bigsqcup x_i)$ is defined. □

Proposition 54 *Let a and b be cpo's. Then $a \times b$, equipped with the componentwise ordering, is a cpo.*

PROOF. Let (z_i) be a partial chain in $a \times b$. Then $(\mathsf{fst}\ z_i)$ and $(\mathsf{snd}\ z_i)$ are partial chains in a and in b, respectively. We thus obtain $\bigsqcup z_i$ as

$$(\bigsqcup(\mathsf{fst}\ z_i), \bigsqcup(\mathsf{snd}\ z_i)).$$

□

Definition 55 We say that a subtype b of a cpo a is a *sub-cpo* of a if the subtype $?b$ of $?a$ is closed under suprema of chains.

Remark 56 It is automatically the case that for a subtype b of a cpo a, $?b$ inherits the bottom element $\bot$ of $?a$, namely as $\lambda \bullet (\bot()\ \mathit{as}\ b)$ (see also Rem. 40).

Proposition 57 *Let a and b be cpo's. Then the type $a \xrightarrow{c} ?b$ of continuous partial functions is a cppo when equipped with the componentwise ordering. The type $a \xrightarrow{c} b$ of continuous total functions is a sub-cpo of $a \xrightarrow{c} ?b$.*

PROOF. Let (f_i) be a partial chain in $a \xrightarrow{c} ?b$. Then $(f_i\ x)$ is a partial chain for all x, so that we obtain $\bigsqcup f_i$ as

$$\left(\lambda x \bullet \bigsqcup(f_i\ x)\right)\ \mathit{res}\ \exists i.\ \mathit{def}\, f_i.$$

It is clear that we can use the same definition for partial chains in $a \xrightarrow{c} b$. The bottom element of $a \xrightarrow{c} ?b$ is $\lambda x \bullet \bot$. □

Proposition 58 *The unit type is a cpo.*

PROOF. The type $?\mathit{Unit} = \mathit{Logical}$ is (internally) even a complete lattice.□

Corollary 59 *If a is a cpo, then $?a$ is a cppo.*

In general, the sum of two cpo's is not again a cpo: as shown in Expl. 50, even $\mathit{Bool} = \mathit{Unit} + \mathit{Unit}$ need not be a (flat) cpo. However, we have

Lemma 60 *In the presence of sum types (Fig. 3), cpo's are stable under sums of partial orders, i.e. the sum $a + b$ of two cpo's a and b is again a cpo when equipped with the sum ordering, iff Bool is a flat cpo.*

PROOF. 'Only if' is trivial. To prove 'if', let a and b be cpo's. Define a function $\mathit{isLeft} : a + b \to \mathit{Bool}$ by

$$\mathit{isLeft}\ z = \mathit{case}\ z\ \mathit{of}\ \mathit{inl}\ x \to \mathit{True} \mid \mathit{inr}\ y \to \mathit{False}.$$

Let s be a partial chain in $a + b$. Then

$$\bigsqcup s_n = \mathit{if}\ \bigsqcup(\mathit{isLeft}\ s_n)\ \mathit{then}\ \mathit{inl}\ \Big(\bigsqcup(\mathit{outl}\ s_n)\Big)\ \mathit{else}\ \mathit{inr}\ \Big(\bigsqcup(\mathit{outr}\ s_n)\Big).$$

□

The precondition of the above lemma has a natural sufficient condition:

Lemma 61 *If Nat is a flat cpo, then so is Bool.*

PROOF. *Bool* is isomorphic to the subtype $\{0, 1\}$ of *Nat*. □

6.2 Domain Datatypes

For use with the concept of recursion laid out in the previous section, HasCasl offers suitable cpo versions of free datatypes. These are declared in otherwise the same syntax as standard free types by means of the keyword **free domain**. Like in the case of free types, the semantics of free domains is defined by means of a fold operation, which however specifies the interpretation of the free domain to be an initial algebra in the above-defined category of internal cpo's — i.e. the fold operation applies only to algebras which are continuous functions on types a of class *Cpo*, returns a continuous function from the initial algebra to a, and is itself continuous. E.g. the specification

var a: *Cpo*
free domain *List* a ::= *nil* | *cons* a (*List* a)

induces an operation

$$\mathit{fold} : c \xrightarrow{c} (a \xrightarrow{c} c \xrightarrow{c} c) \xrightarrow{c} \mathit{List}\ a \xrightarrow{c} c,$$

polymorphic over c : *Cpo*.

The question then arises whether a conservativity result analogous to the one for free types holds for free domains. The answer is positive in case the types of constructor arguments are either types from the same declaration of mutually recursive domains or cpo's from the local environment. This is established by defining a suitable cpo structure on the standard free datatype for the same constructor signature; see the forthcoming extended version of [72] for details. An interesting open problem is whether the result can be extended to types t with non-strict constructors, i.e. constructors with arguments of type $?t$. A

typical example is the type of lazy lists, which can be specified in HasCasl by

var $a : Cpo$
free domain $LList\ a ::= nil \mid cons\ a\ ?(LList\ a)$

where *cons* is a non-strict operation of type $a \xrightarrow{c} ?(LList\ a) \xrightarrow{c} LList$. Intuitively, this type should contain, besides the usual finite lists, finite lazy lists that arise by prefixing elements of a to an undefined lazy list, as well as infinite lists, which come about as suprema of chains of finite lazy lists. It is the subject of ongoing research to formalise this description as a construction in the internal logic of HasCasl, thus showing that the above axiomatic specification is conservative.

6.3 Programming in HasCasl

General recursive function definitions with a cpo-based fixed point semantics may be written in HasCasl as recursive equations in the standard functional programming style, marked by the keyword **program**; these are implicitly translated into the corresponding fixed point terms. One thus obtains essentially a strict functional programming language, in which non-strict functions can be emulated as laid out in Sect. 2.3.

An explicit import of the specification Recursion is not required. A program block is written as a sequence of so-called pattern equations PE_i in the form

program $\{PE_1 \ldots PE_n\}$

Such a program block defines a number of previously declared continuous functions on cpo's. The given types may be partial or total function types; future versions of the HasCasl tools will include termination checks for functions that are declared as total. A pattern equation PE_i has as its left hand side a *pattern* and as its right-hand side an arbitrary term. A pattern is an application of a function being recursively defined to argument terms. In the simplest case, the argument terms are applications of constructors to variables; however, more complex argument patterns including nested patterns, wild cards, tuple patterns, and even Haskell-style as-patterns [58] are also admitted. Variables in patterns need not be explicitly declared; their type is inferred. It is statically checked that all involved types are cpo's; the program block is ill-formed if this check fails. All occurring λ-abstractions, implicit or explicit, are equipped with a downcast to the appropriate continuous function type (so that the user does not have to write these casts explicitly). By consequence,

recursively defined functions are undefined if one of the functions involved in their definition fails to be continuous (sufficient criteria for continuity can be statically checked). Recursive functions on free datatypes can be defined by giving a recursive equation for each constructor. This is coded by means of the case operator; an attempt to use this mechanism for non-free datatypes (which do not have case operators) makes the specification ill-formed. On missing constructor patterns, functions are implicitly undefined; in this case, a warning ('non-exhaustive match') is produced.

As a simple example, Fig. 10 shows an implementation of an interpreter for an abstract imperative core language, where programs are regarded as partial functions on a type $s : Cpo$ of states. The program block is translated into a definition of *eval* as a least fixed point of a continuous functional on the type $Prog \xrightarrow{c} s \xrightarrow{c} ?s$.

spec INTERPRETER = SUMS **then**
    var $s : Cpo$
    free domain $Prog\ s ::= skip \mid basic\ (s \xrightarrow{c} ?\, s)$
        $\mid seq\ (Prog\ s)\ (Prog\ s)$
        $\mid if\ (s \xrightarrow{c} Bool)\ (Prog\ s)\ (Prog\ s)$
        $\mid while\ (s \xrightarrow{c} Bool)\ (Prog\ s)$
    op $eval : (Prog\ s) \xrightarrow{c} s \xrightarrow{c} ?\, s$
    program
        $eval\ skip\ s = s;$
        $eval\ (basic\ f)\ s = f\ s;$
        $eval\ (seq\ p\ q)\ s = eval\ q\ (eval\ p\ s);$
        $eval\ (if\ b\ p\ q)\ s = if\ b\ s\ then\ eval\ p\ s\ else\ eval\ q\ s;$
        $eval\ (while\ b\ p)\ s =$
            $if\ b\ s\ then\ eval\ (while\ b\ p)\ (eval\ p\ s)\ else\ s;$

Fig. 10. Programming an interpreter for a simple abstract language in HASCASL

7 HASCASL in the Heterogeneous Tool Set

Tool support for HASCASL is implemented within the framework of the Bremen heterogeneous tool set Hets [50]. This framework is centred around a *logic graph* in which logics, formalised as institutions (Defn. 23), appear as nodes and logic translations, formalised as *comorphisms* (see below), appear as edges. As a node in the logic graph, HASCASL is equipped with tools for parsing and static analysis. Important translations are an embedding of first order CASL into HASCASL, a connection between HASCASL and the interactive higher order theorem prover Isabelle/HOL, and a mapping of executable

HasCasl specifications into Haskell. We briefly discuss these translations in the following.

Morphisms and Comorphisms of Institutions Recall from Defn. 23 that an institution consists of a category of signatures, equipped with a set-valued sentence functor, a category-valued contravariant model functor, and a satisfaction relation between models and sentences. We briefly recall some standard notions of translations between institutions [23].

Given institutions I and J, an *(institution) comorphism* [23] (also called a *plain map of institutions* [41]) $\mu = (\Phi, \alpha, \beta) : I \to J$ consists of

- a functor Φ from the signature category of I into that of J;
- for each signature Σ in I, a *sentence translation* α_Σ taking Σ-sentences to $\Phi(\Sigma)$-sentences, natural w.r.t. signature morphisms in I; and
- for each signature Σ in I, a *model reduction* functor β_Σ taking $\Phi(\Sigma)$-models to Σ-models, again natural w.r.t. signature morphisms in I,

such that the following *satisfaction condition* holds for all signatures Σ in I, every $\Phi(\Sigma)$-model M, and every Σ-sentence ϕ:

$$M \models_{\Phi(\Sigma)} \alpha_\Sigma \phi \iff \beta_\Sigma M \models_\Sigma \phi.$$

If the model reduction functors β_Σ are surjective on models, then we say that μ is *model-expansive*. In this case, μ admits *borrowing* of entailment systems for basic specifications, i.e. for every signature Σ in I, every set Φ of Σ-sentences, and every Σ-sentence ψ, ψ is a consequence of Φ iff $\alpha_\Sigma(\psi)$ is a consequence of $\alpha_\Sigma[\Phi]$ [11] (where 'only if' holds in general). If μ is even *model-bijective*, i.e. if the β_Σ are bijective on models (but not necessarily on model morphisms), then μ admits borrowing of entailment and refinement also for structured specifications excluding free specifications. If the β_Σ are moreover isomorphisms as functors, then μ admits borrowing of entailment and refinement for structured specifications including free specifications (see [48] for a proof and a detailed explanation of the terminology).

A *theoroidal* comorphism from I to J is a comorphism from I to J^{th}, where J^{th} is obtained from J by replacing the category of signatures with the category of theories (the model functor applied to a theory yields the category of all models of the signature that satisfy the axioms of the theory).

Morphisms of institutions are defined dually to comorphisms: a morphism $\mu = (\Phi, \alpha, \beta) : I \to J$ between institutions I and J consists of a functor Φ from the signature category of I to that of J, sentence translation functions α_Σ from $\Phi(\Sigma)$-sentences to Σ-sentences, and model translation functors β_Σ from

Σ-models to $\Phi(\Sigma)$-models for every signature Σ in I, subject to the obvious satisfaction condition. In heterogeneous specifications, comorphisms appear naturally in translations of structured specifications, while morphisms appear naturally in reductions [49].

Embedding CASL into HASCASL Syntactically, HASCASL is essentially a superset of CASL, so that CASL users can upgrade to HASCASL at liberty. Some subtleties are however attached to the semantic basis of this embedding. It is not possible to define an institution comorphism from CASL into HASCASL: generally, there are no comorphisms from extensional institutions into intensional institutions that work by embedding logical syntax. This is due to the satisfaction condition: e.g., disjunction is extensional in CASL, i.e. for a model M and sentences ϕ, ψ, $M \models \phi \vee \psi$ iff $M \models \phi$ or $M \models \psi$. Now if M is a HASCASL model such that $M \models \phi \vee \psi$ but neither $M \models \phi$ or $M \models \psi$ (such models exist), then there is no possible choice for a reduct $\beta(M)$ of M in the extensional source institution, as by the satisfaction condition $\beta(M)$ would also have to satisfy $\phi \vee \psi$ but none of ϕ and ψ.

The solution to this problem is to connect CASL and HASCASL by a network of morphisms and comorphisms, involving the following modifications of HASCASL.

- The institution HASCASL/FOL of *classical first order* HASCASL is obtained by restricting signatures to first order signatures, in which all operations have types of the form $s_1 \to \ldots \to s_n \to t$ or $s_1 \to \ldots \to s_n \to ?t$, where the s_i are basic types and t is either *Unit* or a basic type, and sentences to formulas not involving λ-abstraction, pairing, quantification over non-basic types, or type variables. Models and satisfaction are inherited from the first level of HASCASL, except that we restrict models to those with a Boolean algebra (rather than just a Heyting algebra) of truth values; moreover, we relax the notion of model morphism to the standard notion of morphism of partial first-order structures (i.e. we require only the weak homomorphism property, which in particular amounts to preservation of definedness for partial functions and preservation of satisfaction for predicates, while higher order homomorphisms may additionally impose reflection in both cases; in particular, this is always the case for higher order homomorphisms of standard models). The most relevant difference between CASL and HASCASL/FOL is that the latter does not have sort generation constraints as a separate type of sentence.
- The institution $\text{HASCASL}^{\text{c,uc}}$ is obtained from HASCASL by restricting models to be classical and to satisfy unique choice. (This is essentially the internal logic of Boolean toposes.)
- The institution $\text{HASCASL}^{\text{ext,uc}}$ is obtained from HASCASL by restricting models to be extensional (Sect. 2.4) and to satisfy unique choice. (This is

essentially the internal logic of well-pointed toposes.)
- The institution HasCasl$^{\text{std}}$ is obtained from HasCasl by restricting models to be standard (Sect. 2.4).

We then have the following morphisms and comorphisms.

- There is an isomorphism Φ from HasCasl/FOL signatures to Casl signatures, and translations α_Σ of Σ-sentences of HasCasl/FOL into $\Phi(\Sigma)$-sentences of Casl; both Φ and α perform only trivial syntactic rearrangements, such as replacing profiles of operators by the corresponding curried function types. In combination with model reductions β_Σ which map a $\Phi(\Sigma)$-model in Casl to the associated standard Σ-model in HasCasl/FOL, we obtain a comorphism (Φ, α, β) : HasCasl/FOL $\to$ Casl (with functoriality of the β_Σ ensured by the relaxation to first-order homomorphisms in HasCasl/FOL). This comorphism fails to be model-expansive and hence does not admit borrowing of entailment; however, the standard (finitary) proof systems for HasCasl/FOL and Casl differ only w.r.t. the absence of sort generation constraints in HasCasl/FOL. Since Φ is an isomorphism, one obtains also a *morphism* $(\Phi^{-1}, \alpha, \beta)$: Casl $\to$ HasCasl/FOL which, intuitively speaking, imports standard models and hides sort generation constraints.
- One has a comorphism Casl $\to$ HasCasl$^{\text{std}}$ which embeds Casl-signatures and standard sentences into HasCasl by acting as an inverse to the corresponding mappings in the comorphism HasCasl/FOL $\to$ Casl above, translates sort generation constraints into induction axioms, and reduces HasCasl models (of essentially first-order signatures) to Casl models by just forgetting higher order structure. Reduction is bijective on models; hence, this comorphism admits borrowing of entailment and refinement for structured specifications not including free specifications.
- Since extensionality implies excluded middle and standard models satisfy unique choice, one has obvious comorphisms HasCasl $\to$ HasCasl$^{\text{c,uc}}$ $\to$ HasCasl$^{\text{ext,uc}}$ $\to$ HasCasl$^{\text{std}}$. These embeddings are isomorphic on signatures and therefore give rise also to morphisms HasCasl$^{\text{std}}$ $\to$ HasCasl$^{\text{ext,uc}}$ $\to$ HasCasl$^{\text{c,uc}}$ $\to$ HasCasl. Additionally, one has a theoroidal comorphism HasCasl$^{\text{c,uc}}$ $\to$ HasCasl which explicitly adds excluded middle and unique choice to each HasCasl signature. None of the translations between HasCasl$^{\text{std}}$, HasCasl$^{\text{ext,uc}}$, and HasCasl$^{\text{c,uc}}$ are model-expansive; nevertheless, the standard (finitary) proof systems for the three logics coincide, so that transitions between HasCasl$^{\text{std}}$ and HasCasl$^{\text{c,uc}}$ are transparent for the user. Model reduction functors are isomorphisms for the comorphism HasCasl$^{\text{c,uc}}$ $\to$ HasCasl, which hence admits borrowing of entailment and refinement for structured specifications including free specifications.
- One has a comorphism HasCasl/FOL $\to$ HasCasl$^{\text{c,uc}}$ which is just a sublogic embedding. In particular, the model reduction functors are iso-

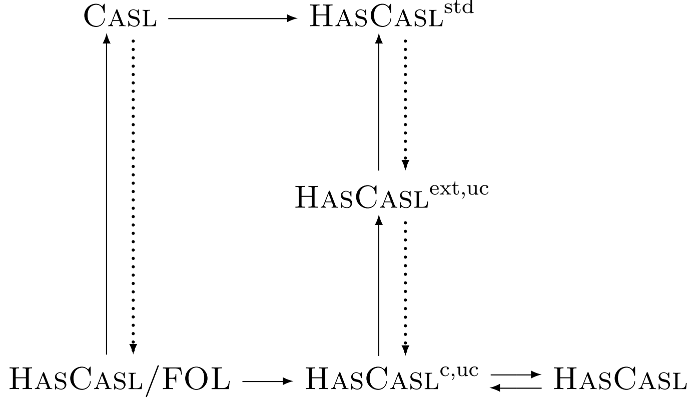

Fig. 11. The CASL-HASCASL network

morphisms, so that this comorphism admits borrowing of entailment and refinement for structured specifications including free specifications.

The ensuing network of morphisms and comorphisms is shown in Fig. 11. Solid lines indicate comorphisms, and dotted lines indicate morphisms.

The comorphism part of the diagram commutes. The sequences of heterogeneous hidings and translations corresponding to the two paths from CASL to HASCASL^{c,uc} in the diagram, however, are distinct, as the path via HASCASL/FOL syntactically hides sort generation constraints, while the path via HASCASL^{std} makes them explicit as induction axioms. The path via HASCASL^{std} will therefore usually be the preferable one; the interest in the path via HASCASL/FOL lies primarily in the fact that it involves only sublogics, rather than fundamental semantic modifications, of HASCASL.

The main question that remains w.r.t. borrowing of entailment systems in the diagram is whether borrowing for structured specifications *including* free specifications is possible between CASL and any of the higher order logics in the diagram. The answer is negative:

Example 62 Consequences of free CASL specifications are not in general preserved by the comorphism into HASCASL^{std}. To see this, consider extending the CASL signature Σ with two sorts s, t and two operations $f, g : s \to t$ to a specification *SP* with an additional sort v, two constants $c, d : v$, and the axiom

- $(\forall x : s \bullet f(x) = g(x)) \Rightarrow c = d$

In CASL, the free specification $SP' = \Sigma$ **then free** *SP* has the consequence $c \neq d$, and hence $f \neq g$ (in particular, SP' is not conservative over Σ): *SP*-models with $f = g$ are never free over their Σ-reducts, since it is always

possible to find a homomorphism from their Σ-reduct into the Σ-reduct of an *SP*-model where $f \neq g$ and $c \neq d$. In HasCasl$^{\mathrm{std}}$, however, the free extension does have models where $f = g$: when f and g are identical functions, then they will remain identical under all higher order homomorphisms. Note that this is a general phenomenon that will occur in any higher order extension of a first order logic that imposes higher order homomorphisms (including intensional higher order logics such as HasCasl, where it may still happen that f and g denote the same element of the function type).

Connecting HasCasl to Isabelle/HOL Proof support for HasCasl is presently based on Isabelle/HOL [54]; the extensible structure of the heterogeneous tool set will however cater for connections to other theorem provers in the future. Indeed the internal logic is to some degree at variance with Isabelle/HOL, as the latter imposes both the law of excluded middle and unique choice (unlike e.g. the Coq proof assistant [14], which therefore presents a future option for alternative proof support). Thus, the comorphism into Isabelle is actually defined on the sublogic HasCasl$^{\mathrm{c,uc}}$ of HasCasl (see above). The mapping from HasCasl$^{\mathrm{c,uc}}$ into Isabelle/HOL codes out subtypes by making subtype injections explicit, and translates partial function types $a \to ?b$ into total function types `'a => 'b option`, where `'b option` is a built-in Isabelle/HOL datatype extending `'b` by an additional element. Finally, the problem that Isabelle does not offer direct support for constructor classes is solved by mapping type constructor variables to loose type constructors, and axioms or implied formulas involving such variables to axioms or proof obligations, respectively (by Thm. 31, this yields a complete proof principle for the extension semantics).

Animating executable HasCasl specifications in Haskell The executable sublogic execHasCasl of HasCasl is defined as follows: a specification belongs to execHasCasl if

- all types it declares are of class *Cpo*,
- all its operation declarations are coherent (i.e. use the keyword **op**) and have types of class *Cpo* (in particular, such types involve only the (total or partial) continuous function type constructors, both of which are mapped to Haskell's function type constructor)
- its only axioms are pattern equations in program blocks. Program blocks are expected to define values of types of class *Cppo*, typically partial continuous functions. Program blocks defining total continuous functions are admitted; they are interpreted as downcasts of fixed points in the partial function types and generate a corresponding termination proof obligation.

It is straightforward to translate EXECHASCASL specifications into Haskell programs, noting that standard λ-abstractions in HASCASL need to be translated into strict λ-abstractions in Haskell, while HASCASL terms of the form $\lambda x :?a \bullet t$ can be translated into standard (non-strict) Haskell-terms `\ x -> t`; otherwise, the only real problems that arise concern the management of name spaces. Thus, HASCASL supports a development methodology where abstract requirement specifications are successively refined into design specifications and eventually executable specifications, which are then automatically translated into Haskell.

8 Monads for Functional-Imperative Programming

We now proceed to establish specification support for imperative constructs, which are embedded into functional programming languages such as Haskell [58] by means of monadic encapsulation of side-effects in the spirit of the seminal paper [46]. We give a brief introduction to the basic concepts of monad-based functional-imperative programming, and then introduce a generic monad-based Hoare logic [74]. (Monad-based Hoare logics discussed in [18,46] are specific for particular types of state monad.)

Intuitively, a monad associates to each type A a type TA of computations of type A; a function with side effects that takes inputs of type A and returns values of type B is, then, just a function of type $A \to TB$. This approach abstracts away from particular notions of computation such as store, non-determinism, non-termination etc.; a surprisingly large amount of reasoning can in fact be carried out independently of the choice of such a notion.

More formally, a monad on a given category $\mathbf{C}$ can be defined as a *Kleisli triple* $(T, \eta, _^*)$, where $T : \mathrm{Ob}\,\mathbf{C} \to \mathrm{Ob}\,\mathbf{C}$ is a function, the *unit* η is a family of morphisms $\eta_A : A \to TA$, and $_^*$ assigns to each morphism $f : A \to TB$ a morphism $f^* : TA \to TB$ such that

$$\eta_A^* = id_{TA}, \qquad f^*\eta_A = f, \qquad \text{and} \qquad g^*f^* = (g^*f)^*.$$

This description is equivalent to the more familiar one via an endofunctor with unit and multiplication [39].

In order to support a language with finitary operations and multi-variable contexts (see below), one needs a further technical requirement: a monad is called *strong* if it is equipped with a natural transformation

$$t_{A,B} : A \times TB \to T(A \times B)$$

called *tensorial strength*, subject to certain coherence conditions (see e.g. [46]);

this is equivalent to enrichment of the monad over **C** (see discussion and references in [46]).

Example 63 [46] Computationally relevant monads on **Set** (since strength is equivalent to enrichment, all monads on **Set** are strong) include

- stateful computations with possible non-termination: $TA = (S \rightarrow ?(A \times S))$, where S is a fixed set of states and $_ \rightarrow ? _$ denotes the partial function type;
- non-determinism: $TA = \mathcal{P}(A)$, where $\mathcal{P}$ denotes the power set functor;
- exceptions: $TA = A + E$, where E is a fixed set of exceptions;
- interactive input: TA is the smallest fixed point of $\gamma \mapsto A + (U \rightarrow \gamma)$, where U is a set of input values.
- non-deterministic stateful computations: $TA = (S \rightarrow \mathcal{P}(A \times S))$, where, again, S is a fixed set of states;
- continuations: $TA = (A \rightarrow R) \rightarrow R$, where R is a type of results.

In order to accommodate binding also of programs $A \rightarrow ?TB$ with intrinsic non-termination, we use the specification of monads already shown in Fig. 5. It is slightly modified w.r.t. the standard laws for monads, the main subtlety arising from partiality being the treatment of the first unit law [75]. The notation is (almost) identical to the one used in Haskell, i.e. the unit is denoted by *ret*, and the binding operator $_ \gg= _$ denotes, in Kleisli triple notation, the function $(x, f) \mapsto f^*(x)$. There is a built-in sugaring for the binding operation in the form of a Haskell-style *do*-notation: for monadic expressions p and q,

$$do\ x \leftarrow p;\ q$$

abbreviates $p \gg= \lambda x \bullet q$. (This is essentially the same as Moggi's *let*-notation [46].) The intuition behind this notation is that the computations p and q are performed sequentially, with the result of p being bound to x and passed on to q (an expression which may contain the variable x).

In the do-notation, the axioms of Fig. 5 take the following shape: Binding is *associative*, i.e. one has

$$do\ y \leftarrow (do\ x \leftarrow p;\ q);\ r = do\ x \leftarrow p;\ do\ y \leftarrow q;\ r$$

if r does not contain x. Moreover, we have *unit laws* stating that

$$\begin{aligned}(do\ x \leftarrow ret\ a;\ p) &= p[x/a], \text{ whenever } p[x/a] \text{ is defined,}\\ (do\ y \leftarrow q; x \leftarrow ret\ a;\ p) &= do\ y \leftarrow q;\ p[x/a], \text{ and}\\ (do\ x \leftarrow p;\ ret\ x) &= p.\end{aligned}$$

Thanks to associativity, one may safely denote nested *do* expressions like $do\ x \leftarrow p;\ do\ y \leftarrow q;\ \dots$ by $do\ x \leftarrow p; y \leftarrow q;\ \dots$. Repeated nestings such as $do\ x_1 \leftarrow p_1, \dots, x_n \leftarrow p_n;\ q$ are somewhat inaccurately denoted in

the form $do\ \bar{x} \leftarrow \bar{p};\ q$. Term fragments of the form $\bar{x} \leftarrow \bar{p}$ are called *program sequences*. Bound variables x_i that are not used later may be omitted from the notation. Terms are generally formed in a context $\Gamma = (x_1 : s_1, \ldots, x_n : s_n)$ of variables with assigned types. Following [46], we shall refer to this notation and the associated calculus as the *computational meta-language*.

As an example of an instance of the class *Monad*, a specification of the state monad is shown in Fig. 12. Note that it is only thanks to the treatment of partial functions in the specification of monads that the state monad is really an instance of *Monad*, since stricter versions of the first monad law fail to hold for the state monad (see [75] for a more detailed discussion). Monads specified in HasCasl in the style of Fig. 12 are automatically strong, because the operations of the monad are internalised as functions (recall that strength is equivalent to enrichment).

```
spec  STATE = MONAD then
      var      state : Type
      type instance ST state : Monad
      vars     a, b : Type
      type     ST state a := state →? (a × state)
      var      x : a;  p : ST state a;  q : a →? ST state b
               • (ret x) : ST state a = λ s : state • (x, s)
               • (p >>= q) = λ s1 : state • let (z, s2) = p s1 in q z s2
```

Fig. 12. Specification of the state monad

On top of a monad, one can generically define control structures such as if-then-else or loops. The if-then-else construct is defined by

$$if\ b\ then\ p\ else\ q = do\ a \leftarrow b;\ if\ a\ then\ p\ else\ q$$

for $b : T\,Bool$ and $p, q : TA$, where the stateless if-then-else construct on the right hand side is the one defined in Fig. 3. Loops require general recursion on function spaces between flat cpo's. Since in the absence of unique choice, not all types need be flat cpo's when equipped with the equality ordering (see Sec. 6), one thus needs to restrict to monads that preserve flat cpo's (under unique choice, this condition is void). This is an example of a constructor subclass; the corresponding specification of *flat cpo monads* is shown in Fig. 13. The specification declares the class *FlatCpoMonad* to be a subclass of both *FlatCpo* $\rightarrow$ *FlatCpo* and *Monad*, i.e. a flat cpo monad is a monad which restricts to the class *FlatCpo*; moreover we require that the binding operation is continuous on flat cpo's. Note that for flat cpo's a, b, the continuous function types $a \xrightarrow{c} b$ and $a \xrightarrow{c}?b$ coincide with the respective function types $a \rightarrow b$ and $a \rightarrow?b$, so that there is no need to explicitly specify continuity of the return operation. Most relevant computational monads including the ones in Expl. 63 above are instances of this subclass (even without unique choice).

spec FLATCPOMONAD = RECURSION **and** MONAD **then**
class *FlatCpoMonad* $<$ *FlatCpo* $\to$ *FlatCpo*
class *FlatCpoMonad* $<$ *Monad*
var m: *FlatCpoMonad*; a, b: *FlatCpo*
op $_ \gg= _ : m\,a \times (a \stackrel{c}{\longrightarrow}?\,m\,b) \stackrel{c}{\longrightarrow}?\,m\,b$;

Fig. 13. The constructor subclass of flat cpo monads

As an example of a loop construct we introduce an iteration construct which generalises the while loop by extending it with a default return value (the while loop as programmed e.g. in the Haskell prelude returns only a unit value) which is fed through the iteration. The specification of the iteration construct is shown in Fig. 14. Note that the while loop is just iteration with a dummy return value.

spec ITERATION = SUMS **and** FLATCPOMONAD **then**

vars m: *FlatCpoMonad*; a: *FlatCpo*
ops $iter : (a \stackrel{c}{\longrightarrow}?\,m Bool) \stackrel{c}{\longrightarrow} (a \stackrel{c}{\longrightarrow}?\,m\ a) \stackrel{c}{\longrightarrow} a \stackrel{c}{\longrightarrow}?\,m\ a$;
$while : m\ Bool \stackrel{c}{\longrightarrow} m\ Unit \stackrel{c}{\longrightarrow}?\,m\ Unit$
program
iter test f x = do b $\leftarrow$ *test x*
if b then
do y $\leftarrow$ *f x*; *iter test f y*
else ret x
while b p = iter $(\lambda \bullet!\ b)\ (\lambda \bullet p)\ ()$

Fig. 14. The iteration control structure

Remark 64 The iteration construct may more generally be defined by recursion on general cpo's. This requires a specification of monads on the category of cpo's and continuous functions, in perfect analogy to the specification given in Fig. 5 which defines the class of monads on the category of types and functions. The most convenient way to express this is to parametrise the specification of monads over the base category of the monad, i.e. a class equipped with subtypes of the function types representing the morphisms. We avoid such a parametrised specification of monads purely in the interest of readability.

9 Generic Purity and Global Evaluation

In preparation for the formulation of the monad-based Hoare calculus, we now summarise material from [75] on generic notions of purity (previously called side-effect freeness), to be required of stateful formulas appearing as pre- and postconditions, and *global evaluation formulas* (called *global dynamic judge-*

ments in [75]). Informally, pure programs are those that have the following properties:

(1) discardability: if their result is not used, then pure computations can be left out from a sequence of computation steps without changing its behaviour;
(2) determinism (copyability): when executed repeatedly, pure programs always return the same value;
(3) interchangeability: pure programs can be interchanged with each other, that is, the order does not matter.

Purity is generally a much weaker property than statelessness, which means that there is no interaction with 'state', or generally the monad, at all.

We fix the notation for monads introduced in the previous section (T, η etc.) throughout the remaining development.

Definition 65 [19,79] A program p is called *stateless* if it factors through *ret*, i.e. if it is just a value inserted into the monad. A program p is called *discardable* if

$$(do\ y \leftarrow p;\ \ ret\ *) = ret\ *,$$

where $*$ is the unique element of the unit type. A program p is called *copyable* if

$$(do\ x \leftarrow p; y \leftarrow p;\ \ ret\ (x, y)) = do\ x \leftarrow p;\ \ ret\ (x, x)$$

for $x \notin FV(p)$, where $FV(p)$ denotes the set of free variables in p. Moreover, programs p, q *commute* if

$$(do\ x \leftarrow p; y \leftarrow q;\ \ ret\ (x, y)) = do\ y \leftarrow q; x \leftarrow p;\ \ ret\ (x, y)$$

for $x \notin FV(q)$, $y \notin FV(p)$.

Proposition and Definition 66 *[75] Let p be discardable and copyable. Then p commutes with all discardable copyable programs iff p commutes with all discardable copyable Logical-valued programs. In this case, p is called* pure. *The subtype of TA formed by the pure computations will be denoted by PA throughout.*

For details on the relation between the various notions above, see [19,75]. Here, we need mainly the notion of purity. Stateless programs are pure, but not conversely. For example, in the state monad, statelessness means that the program neither changes nor reads the state (p is stateless iff p *exists* in the sense of [46]). Contrastingly, we have

Example 67 A program p is pure

- in the state monad iff p terminates and does not change the state (p may

however read the state);
- in the non-determinism monad iff p has a unique outcome;
- in the exception monad iff p terminates normally;
- in the interactive input monad iff p never reads any input;
- in the non-deterministic state monad iff p does not change the state and always has a unique outcome;
- in the continuation monad (over **Set**) iff p is stateless.

The definition of the semantics of the Hoare logic is based on *global evaluation formulas* $[\bar{x} \leftarrow \bar{p}]\,\phi$, where $\bar{x} \leftarrow \bar{p}$ is a program sequence and $\phi : \mathit{Logical}$ is a formula which may contain $\bar{x}$. Intuitively, $[\bar{x} \leftarrow \bar{p}]\,\phi$ states that ϕ holds for the result values $\bar{x}$ after execution of $\bar{x} \leftarrow \bar{p}$ from any initial state. Formally, $[\bar{x} \leftarrow \bar{p}]\,\phi$ abbreviates

$$(\mathit{do}\ \bar{x} \leftarrow \bar{p};\ \mathit{ret}\ (\bar{x}, \phi)) = \mathit{do}\ \bar{x} \leftarrow \bar{p};\ \mathit{ret}\ (\bar{x}, \top)$$

(a strong equation). The degenerate case $[\,]\,\phi$ is (by injectivity of *ret* as specified in Fig. 5) equivalent to ϕ; we shall silently identify the two formulas.

Remark 68 The above semantics of global evaluation formulas is close to Moggi's global semantics of evaluation logic [47] (but not at all to the original local semantics as defined in [62], which is related instead to the monad-based dynamic logic of [75,52]).

Example 69 In the monads of Expl. 63, satisfaction of $[x \leftarrow p]\,\phi$, where $p : TA$, amounts to the following (we freely omit semantic brackets from the notation):

- *states*: terminating execution of p from any initial state yields a result value x satisfying ϕ;
- *non-determinism*: all values x in $p \in \mathcal{P}(A)$ satisfy ϕ;
- *exceptions*: if p terminates normally, then its result value x satisfies ϕ;
- *interactive input*: the value x eventually produced by p after some combination of inputs always satisfies ϕ;
- *non-deterministic state monad*: all possible result values x obtained by execution of p from any initial state satisfy ϕ;
- *continuations*: for $k : A \to R$, $p\ k$ depends only on the restriction of k to the set of values $x : A$ satisfying ϕ.

Figure 15 shows a number of proof rules for global evaluation formulas. In the present setting, this should be regarded as a collection of lemmas rather than as a formally delimited calculus; in particular, we shall apply the rules in proofs using the full power of the ambient higher order logic. A slightly different calculus for a clearly separated definition of *global evaluation logic* is given in [24]. Double lines indicate that a rule works in both directions. Recall that $FV(p)$ denotes the set of free variables of p. The rules (pre) and (wk) use

explicit quantification to enforce the usual variable condition stating that certain variables do not occur freely in assumptions. We have formulated specific rules for pure terms; some of these hold more generally e.g. for discardable or copyable terms, respectively, but we will not need this added generality. Soundness of the rules, and derivability of the rules marked as such, has been established in [75] (Figs. 3 and 4 and Prop. 4.29), except rule (ins) which is derived using (app), (pre), and (comm). We will refer to proofs using only the rules ($\wedge$I) and (wk) as *propositional reasoning.*

Basic rules

$$(\wedge\mathbf{I})\ \frac{[\bar{x} \leftarrow \bar{p}]\,\phi \qquad [\bar{x} \leftarrow \bar{p}]\,\xi}{[\bar{x} \leftarrow \bar{p}]\,(\phi \wedge \xi)} \qquad (\mathbf{wk})\ \frac{\forall \bar{x}.\,\phi \Rightarrow \xi \qquad [\bar{x} \leftarrow \bar{p}]\,\phi}{[\bar{x} \leftarrow \bar{p}]\,\xi} \qquad (\mathbf{eq})\ \frac{[\bar{x} \leftarrow \bar{p}]\,q_1 = q_2 \qquad [\bar{x} \leftarrow \bar{p}; y \leftarrow q_1; \bar{z} \leftarrow \bar{r}]\,\phi}{[\bar{x} \leftarrow \bar{p}; y \leftarrow q_2; \bar{z} \leftarrow \bar{r}]\,\phi}$$

$$(\mathbf{app})\ \frac{[\bar{x} \leftarrow \bar{p}]\,\phi \qquad y \notin FV(\phi)}{[\bar{x} \leftarrow \bar{p}; y \leftarrow q]\,\phi} \qquad (\mathbf{pre})\ \frac{\forall x.\,[\bar{y} \leftarrow \bar{q}]\,\phi}{[x \leftarrow p; \bar{y} \leftarrow \bar{q}]\,\phi}$$

$$(\eta)\ \frac{[\bar{x} \leftarrow \bar{p}; y \leftarrow \mathit{ret}\ a; \bar{z} \leftarrow \bar{q}]\,\phi}{\overline{\overline{[\bar{x} \leftarrow \bar{p}; \bar{z} \leftarrow \bar{q}[a/y]]\,\phi[a/y]}}} \qquad (\mathbf{ctr})\ \frac{[\ldots; x \leftarrow p; y \leftarrow q; \bar{z} \leftarrow \bar{r}]\,\phi \qquad x \notin FV(\phi) \cup FV(\bar{r})}{[\ldots; y \leftarrow (\mathit{do}\ x \leftarrow p;\ q); \bar{z} \leftarrow \bar{r}]\,\phi}$$

. . . .

Rules for pure terms

$$(\mathbf{dis}_0)\ \frac{[\bar{x} \leftarrow \bar{p}; q]\,\phi \qquad q \text{ pure}}{[\bar{x} \leftarrow \bar{p}]\,\phi} \qquad (\mathbf{comm})\ \frac{[\bar{x} \leftarrow \bar{p}; y \leftarrow q; z \leftarrow r; \bar{w} \leftarrow \bar{s}]\,\phi \qquad q, r \text{ pure} \qquad y \notin FV(r), z \notin FV(q)}{[\bar{x} \leftarrow \bar{p}; z \leftarrow r; y \leftarrow q; \bar{w} \leftarrow \bar{s}]\,\phi}$$

$$(\mathbf{copy})\ \frac{[\bar{x} \leftarrow \bar{p}; y \leftarrow q; z \leftarrow q; \bar{w} \leftarrow \bar{r}]\,\phi}{\overline{\overline{[\bar{x} \leftarrow \bar{p}; y \leftarrow q; \bar{w} \leftarrow \bar{r}[y/z]]\,\phi[y/z]}}}\ (q \text{ pure}, y \notin FV(q))$$

. . . .

Derived rules

$$(\mathbf{tau})\ \frac{\forall \bar{x}.\,\phi}{[\bar{x} \leftarrow \bar{p}]\,\phi} \qquad (\mathbf{rp})\ \frac{\forall \bar{x}.\,q_1 = q_2 \qquad [\bar{x} \leftarrow \bar{p}; y \leftarrow q_1; \bar{z} \leftarrow \bar{r}]\,\phi}{[\bar{x} \leftarrow \bar{p}; y \leftarrow q_2; \bar{z} \leftarrow \bar{r}]\,\phi}$$

$$(\mathbf{dis})\ \frac{[\bar{x} \leftarrow \bar{p}; q; \bar{z} \leftarrow \bar{r}]\,\phi \qquad q \text{ pure}}{[\bar{x} \leftarrow \bar{p}; \bar{z} \leftarrow \bar{r}]\,\phi} \qquad (\mathbf{ins})\ \frac{[\bar{x} \leftarrow \bar{p}; \bar{z} \leftarrow \bar{r}]\,\phi \qquad q, \bar{r} \text{ pure or } q, \bar{p} \text{ pure}}{[\bar{x} \leftarrow \bar{p}; q; \bar{z} \leftarrow \bar{r}]\,\phi}$$

Fig. 15. Proof rules for global evaluation formulas

Convention 70 Pure terms can be handled notationally in a more relaxed way, as it is immaterial how often and in which order they are evaluated as long as no other programs interfere. We thus allow pure programs of type PA to occur in places where a term of type A is expected in programs and formulas. More precisely, if $\bar{x} = (x_1, \ldots, x_n)$ is a list of variables of types $A_1, \ldots, A_n$ and q is a program, then the program $q[\bar{p}/\bar{x}]$ obtained by substituting terms $p_i : PA_i$ for the x_i is defined as $do\ \bar{x} \leftarrow \bar{p};\ q$, with well-definedness guaranteed by purity (see [75] for details). Similarly, $[\bar{y} \leftarrow \bar{q}]\, \phi[\bar{p}/\bar{x}]$ abbreviates $[\bar{y} \leftarrow \bar{q}; \bar{x} \leftarrow \bar{p}]\, \phi$. Note that this includes the case that $\bar{y} \leftarrow \bar{q}$ is the empty sequence. Since we further identify $[]\, \phi$ with ϕ, e.g. $\phi \Rightarrow \psi$ abbreviates $[a \leftarrow \phi; b \leftarrow \psi]\, (a \Rightarrow b)$ for $\phi, \psi : P\mathit{Logical}$. Ambiguities may arise from polymorphic predicates and operations such as equality, e.g. in the equation $p = q$, with $p, q : PA$. In such cases, we will disambiguate formulas by explicit type annotations where necessary; e.g., $p =_A q$ abbreviates $[x \leftarrow p; y \leftarrow q]\, x = y$, while $p =_{PA} q$ is just equality of computations. A single warning is required: rule (app) of Fig. 15 is sound only if the formula ϕ is really stateless.

10 The generic Hoare calculus

We now proceed to describe the generic monad-based Hoare-calculus.

Definition 71 A *Hoare triple*, written $\{\phi\}\ \bar{x} \leftarrow \bar{p}\ \{\psi\}$, consists of a program sequence $\bar{x} \leftarrow \bar{p}$, a precondition $\phi : T\mathit{Logical}$, and a postcondition $\psi : T\mathit{Logical}$ (which may contain $\bar{x}$), where ϕ and ψ are pure. This abbreviates the global evaluation formula

$$[a \leftarrow \phi; \bar{x} \leftarrow \bar{p}; b \leftarrow \psi]\, (a \Rightarrow b)$$

with fresh variables $a, b : \mathit{Logical}$.

The fact that Hoare triples as just defined mention program *sequences* (rather than just programs) reflects the need to actually reason about results of computations, including intermediate results, as opposed to just about state changes as in the traditional case.

Example 72 A Hoare triple $\{\phi\}\ x \leftarrow p\ \{\psi\}$ holds

- in the state monad iff, whenever ϕ holds in a state s and p terminates in state s' with result x when executed in state s, then ψ holds for x in the state s';
- in the non-determinism monad iff, whenever ϕ is true, then ψ holds for all possible results x of p;
- in the exception monad iff, whenever ϕ holds and p terminates normally, returning x, then ψ holds for x;

- in the interactive input monad iff, whenever ϕ holds and p returns x after reading some sequence of inputs, then ψ holds for x.
- in the non-deterministic state monad iff, whenever ϕ holds in a state s, and p possibly terminates in a state s' with result x, then ψ holds for x in s'.

A set of monad-independent Hoare rules is shown in Fig. 16. There is presently no claim that the calculus is (relatively) complete; the main purpose of the given rule set is to show that large parts of the program verification process can be performed independently of the underlying notions of side-effect. The rules (pure), (wk), (disj), and (conj) apply the notation introduced in Convention 70. In particular, $\phi \Rightarrow \psi$ has the same decoding as the Hoare triple $\{\phi\}\ \{\psi\}$, so that (wk) is actually a special case of the sequential rule (seq). Due to discardability, the decoding of $\phi \Rightarrow \psi$ can be simplified to

$$(do\ a \leftarrow \phi, b \leftarrow \psi;\ ret\ (a \Rightarrow b)) = ret\ \top.$$

In the pre- and postconditions, Boolean values b are implicitly converted to *Logical* as $b = true$, and *formulas of type Logical are implicitly cast to* $P Logical$ via *ret* when needed (used in Fig. 16 only for the formula $\bot : Logical$). Square brackets indicate reasoning with local assumptions, discharged by application of the rule; this occurs only in rule (Y). Universal quantifiers on Hoare triples in premises (rules (seq), (wk), (Y)) are, as already in Fig. 15, just a short way of expressing the variable condition. An exception is the universal quantifier on the assumption in (Y), which means that the derivation may use arbitrary instances of the assumption. Arguments in the calculus using only the rules (*Logical*), ($\bot$), (wk), (conj), and (disj) are referred to as *propositional reasoning.*

The rule (pure) applies in particular to stateless programs $p = ret\ a$, for which the precondition simplifies to $\phi[a/x]$ (see Convention 70). Although the classical Hoare calculus does not require the usual introduction and elimination rules for logical connectives, such rules are sometimes convenient (see the example below); we have included introduction rules for conjunction and disjunction. One typical Hoare rule that is missing here is the assignment rule; this rule only makes sense in a more specialised context where some sort of store is present (the rule (pure) should not be confused with the assignment rule — it refers to the monadic binding mechanism and not to assignment to store locations). An example of an extension of the calculus by specialised rules for a particular monad is presented below. Rule (Y) refers to the fixed-point operator Y (Sect. 6); this rule applies only to flat cpo monads (Sec. 8). Application of the Y operator to F requires implicitly that F has the continuous function type $(A \xrightarrow{c} ?TB) \xrightarrow{c} (A \xrightarrow{c} ?TB)$ for flat cpo's A, B. From (Y), one derives e.g. a rule for the iteration construct from Sect. 8:

$$(\bot)\ \frac{}{\{\bot\}\ p\ \{\varphi\}} \qquad \textbf{(pure)}\ \frac{p \text{ pure}}{\{\phi[p/x]\}\ x \leftarrow p\ \{\phi\}}\ (x \notin FV(p))$$

$$(\mathit{Logical})\ \frac{}{\{\mathit{ret}\ \phi\}\ p\ \{\mathit{ret}\ \phi\}} \qquad (\eta)\ \frac{\{\phi\}\ x \leftarrow p; y \leftarrow \mathit{ret}\ a; z \leftarrow q\ \{\psi\}}{\{\phi\}\ x \leftarrow p; z \leftarrow q[a/y]\ \{\psi[a/y]\}}$$

$$\textbf{(seq)}\ \frac{\{\phi\}\ \bar{x} \leftarrow \bar{p}\ \{\psi\} \quad \forall \bar{x} \bullet \{\psi\}\ \bar{y} \leftarrow \bar{q}\ \{\chi\}}{\{\phi\}\ \bar{x} \leftarrow \bar{p}; \bar{y} \leftarrow \bar{q}\ \{\chi\}} \qquad \textbf{(wk)}\ \frac{\{\phi\}\ \bar{x} \leftarrow \bar{p}\ \{\psi\} \quad \phi' \Rightarrow \phi \quad \forall \bar{x} \bullet \psi \Rightarrow \psi'}{\{\phi'\}\ \bar{x} \leftarrow \bar{p}\ \{\psi'\}}$$

$$\textbf{(ctr)}\ \frac{\{\phi\}\ \ldots; x \leftarrow p; y \leftarrow q; \bar{z} \leftarrow \bar{r}\ \{\psi\}}{\{\phi\}\ \ldots; y \leftarrow (\mathit{do}\ x \leftarrow p;\ q); \bar{z} \leftarrow \bar{r}\ \{\psi\}}\ (x \notin FV(\bar{r}, \psi))$$

$$\textbf{(if)}\ \frac{\{\phi\}\ a \leftarrow b\ \{\mathit{if}\ a\ \mathit{then}\ \psi\ \mathit{else}\ \xi\} \quad \{\psi\}\ x \leftarrow p\ \{\chi\} \quad \{\xi\}\ x \leftarrow q\ \{\chi\}}{\{\phi\}\ x \leftarrow \mathit{if}\ b\ \mathit{then}\ p\ \mathit{else}\ q\ \{\chi\}}\ (a \notin FV(\psi, \xi))$$

$$\textbf{(conj)}\ \frac{\{\phi\}\ \bar{x} \leftarrow \bar{p}\ \{\psi\} \quad \{\phi\}\ \bar{x} \leftarrow \bar{p}\ \{\chi\}}{\{\phi\}\ \bar{x} \leftarrow \bar{p}\ \{\psi \wedge \chi\}} \qquad \textbf{(disj)}\ \frac{\{\phi\}\ \bar{y} \leftarrow \bar{q}\ \{\chi\} \quad \{\psi\}\ \bar{y} \leftarrow \bar{q}\ \{\chi\}}{\{\phi \vee \psi\}\ \bar{y} \leftarrow \bar{q}\ \{\chi\}}$$

$$\begin{array}{c} \left[\forall y \bullet \{\phi\}\ x \leftarrow p\ y\ \{\psi\}\right] \\ \vdots \end{array}$$
$$\textbf{(Y)}\ \frac{\forall y \bullet \{\phi\}\ x \leftarrow F\ p\ y\ \{\psi\}}{\{\phi\}\ x \leftarrow Y\ F\ y\ \{\psi\}}$$

Fig. 16. The generic Hoare calculus (rule (Y) applies only to flat cpo monads)

Proposition 73 *Given the definition of the iteration construct, the rule*

$$\textbf{(iter)}\ \frac{\forall x \bullet \{\phi\}\ a \leftarrow b\ x\ \{\mathit{if}\ a\ \mathit{then}\ \psi\ \mathit{else}\ \xi\} \quad \forall x \bullet \{\psi\}\ y \leftarrow p\ x\ \{\phi[y/x]\}}{\{\phi[e/x]\}\ y \leftarrow \mathit{iter}\ b\ p\ e\ \{\xi[y/x]\}}\ (y \notin FV(\phi, \xi))$$

is derivable in the generic Hoare calculus.

PROOF. Let F be the functional from the definition of the $\mathit{iter}\ b\ p$, i.e.

$$F\ f\ x = \mathit{if}\ b\ \mathit{then}\ (\mathit{do}\ z \leftarrow p\ x;\ f\ z)\ \mathit{else}\ \mathit{ret}\ x.$$

Assume $\forall x \bullet \{\phi\}\ y \leftarrow f\ x\ \{\xi[y/x]\}$. By rule (Y), it suffices to derive

$$\{\phi\}\ y \leftarrow F\ f\ x\ \{\xi[y/x]\}.$$

By rule (if) and the first premise, this reduces to

$$\{\psi\}\ y \leftarrow (do\ z \leftarrow p\ x;\ f\ z)\ \{\xi[y/x]\} \quad \text{and}$$
$$\{\xi\}\ y \leftarrow ret\ x\ \{\xi[y/x]\}.$$

The second goal is discharged immediately by applying (pure), as $\xi[y/x][x/y] = \xi$ due to $y \notin FV(\xi)$. By the assumption and rules (ctr) and (seq), the first goal reduces to

$$\{\psi\}\ z \leftarrow p\ x\ \{\phi[z/x]\},$$

i.e. to the second premise. □

The rules for if-then-else and iteration have been formulated so as to allow side-effecting expressions as conditions. If the condition b is pure, then one derives from the given rules and rule (pure) the usual if rule and a rule for iter corresponding to the standard while rule:

$$\frac{\{\phi \wedge b\}\ x \leftarrow p\ \{\psi\} \quad \{\phi \wedge \neg b\}\ x \leftarrow q\ \{\psi\}}{\{\phi\}\ x \leftarrow if\ b\ then\ p\ else\ q\ \{\psi\}} \qquad \frac{\{\phi \wedge (b\ x)\}\ y \leftarrow p\ x\ \{\phi[y/x]\}}{\{\phi[e/x]\}\ y \leftarrow iter\ b\ p\ e\ \{\phi[y/x] \wedge \neg(b\ y)\}}.$$

The latter rule specialises to the usual while rule

$$\textbf{(while)}\ \frac{\{\phi \wedge b\}\ p\ \{\phi\}}{\{\phi\}\ while\ b\ p\ \{\phi \wedge \neg b\}}.$$

The rules of the calculus are sound for arbitrary (flat cpo) monads:

Theorem 74 *If a Hoare triple is derivable in a (flat cpo) monad by the rules of Fig. 16 excluding (including) rule (Y), then its decoding is derivable in the internal language.*

PROOF. We prove each rule as a lemma in the internal language, using the proof rules of Fig. 15:

(*Logical*), ($\bot$), (η): Straightforward from rules (tau) and (η) of Fig. 15.

(pure): Renaming the bound occurrence of x to a fresh variable y and taking into account Convention 70, we decode the conclusion to

$$[x \leftarrow p; a \leftarrow \phi; y \leftarrow p; b \leftarrow \phi[y/x]]\ a \Rightarrow b$$

where a, b are fresh. Since all involved terms are pure, this reduces to

$$[x \leftarrow p; a \leftarrow \phi]\, a \Rightarrow a$$

by rules (comm) and (copy) of Fig. 15. The latter formula is immediate by rule (tau) of Fig. 15.

(seq): By rules (app) and (pre) of Fig. 15, the premises imply

$$[a \leftarrow \phi; x \leftarrow \bar{p}; b \leftarrow \psi; \bar{y} \leftarrow \bar{q}; c \leftarrow \chi]\, a \Rightarrow b \quad \text{and}$$
$$[a \leftarrow \phi; x \leftarrow \bar{p}; b \leftarrow \psi; \bar{y} \leftarrow \bar{q}; c \leftarrow \chi]\, b \Rightarrow c.$$

By propositional reasoning, we obtain

$$[a \leftarrow \phi; x \leftarrow \bar{p}; \psi; y \leftarrow q; c \leftarrow \chi]\, a \Rightarrow c.$$

The conclusion then follows by the rule (dis) of Fig. 15.

(wk): As indicated above, this is a special case of (two applications of) (seq).

(ctr): Immediate by rule (ctr) of Fig. 15.

(if): Since *if b then p else q* is just an abbreviation for *do* $a \leftarrow b$; *if a then p else q*, the conclusion reduces by rules (seq) and (ctr) and the first premise to

$$\{\mathit{if}\ a\ \mathit{then}\ \psi\ \mathit{else}\ \xi\}\ x \leftarrow \mathit{if}\ a\ \mathit{then}\ p\ \mathit{else}\ q\ \{\chi\}$$

for $a : \mathit{Bool}$. We can then perform a case distinction over a. If $a = \mathit{True}$, then the above formula is equivalent (by rule (rp) of Fig. 15) to

$$\{\psi\}\ x \leftarrow p\ \{\chi\},$$

i.e. the second premise. The case $a = \mathit{False}$ is analogous.

(conj): By rule (ins) of Fig. 15, we obtain from the premises

$$[a \leftarrow \phi; \bar{x} \leftarrow \bar{p}; b \leftarrow \psi; c \leftarrow \chi]\, a \Rightarrow b \quad \text{and}$$
$$[a \leftarrow \phi; \bar{x} \leftarrow \bar{p}; b \leftarrow \psi; c \leftarrow \chi]\, a \Rightarrow c.$$

By propositional reasoning, this implies

$$[a \leftarrow \phi; \bar{x} \leftarrow \bar{p}; b \leftarrow \psi; c \leftarrow \chi]\, a \Rightarrow b \wedge c.$$

By rule (η) of Fig. 70, we obtain

$$[a \leftarrow \phi; \bar{x} \leftarrow \bar{p}; b \leftarrow \psi; c \leftarrow \chi; d \leftarrow \mathit{ret}\ (b \wedge c)]\, a \Rightarrow d,$$

which is precisely the decoding of the conclusion, taking into account Convention 70.

(disj): Analogous to (conj).

(Y): Let $F : (A \xrightarrow{c} ?TB) \xrightarrow{c} (A \xrightarrow{c} ?TB)$ for flat cpo's A, B. As the bottom element $\bot$ of $A \xrightarrow{c} ?TB$ satisfies $\forall y \bullet \{\phi\}\ x \leftarrow \bot\ y\ \{\psi\}$, correctness of the rule follows by fixed point induction if the predicate $\lambda z : A \xrightarrow{c} ?TB \bullet \forall y \bullet \{\phi\}\ x \leftarrow z\ y\ \{\psi\}$ is admissible, i.e. closed under suprema of total chains. This is easily established in the internal logic, noting that $(\bigsqcup f_i)\ x = \bigsqcup (f_i\ x)$, that Hoare triples decode into equations between *do*-terms, and that binding is continuous in flat cpo monads. □

Completeness of the calculus for the class of all (flat cpo) monads is the subject of ongoing research. It is clear that completeness of the calculus over a specific monad can only be expected in combination with suitable monad-specific rules; e.g., the calculus becomes the usual (complete) Hoare calculus when extended with an assignment rule specific to the store monad. In this sense, the calculus may be regarded as a generic framework for computational deduction systems.

11 Example: Reasoning about dynamic references

We now apply the general machinery developed so far to the (slightly extended) domain of the classical Hoare calculus, namely states consisting of creatable and destructively updatable references (note that this is just one example of a state monad), later to be extended by non-determinism.

The specification of reference monads is shown in Fig. 17. It uses a type constructor *Ref*, where *Ref* a is the set of references to values of type a. All reference types are made subtypes of a fixed type *Loc* of locations, which allows comparing references of different type. Nothing is said a priori on whether references of different type must be distinct as locations. In dynamic reference monads according to the specification DynamicReference, however, distinctness of references may be inferred in all relevant cases from their separate creation. The monad comes with operations for reading from and writing to references (besides the usual monad operations). The read operation * is pure, which is expressed in the axiom *pure-read* using a built-in predicate *pure*; by Convention 70, this allows using the read operation in places where values are expected. Note the difference between $r = s$ (equality of references, a stateless formula) and $*r = *s$ (equality of contents, a stateful formula).

The axiomatisation provides all that is really necessary in order to reason about references, i.e. one does not need to rely on a particular implementation. Axiom *read-write* says that after writing to a reference, we can read the

spec REFERENCE = FLATCPOMONAD **then**
type *Loc*
var $a\colon FlatCpo$
types $R\colon FlatCpoMonad$;
$Ref\, a < Loc$; $Ref\ a\colon FlatCpo$
ops $*\colon Ref\ a \xrightarrow{c} R\ a$;
$_ := _\colon (Ref\ a) \times a \xrightarrow{c} R\ Unit$
var $x\colon a;\ y\colon b;\ r\colon Ref\ a;\ s\colon Ref\ b$
• $pure\ (*r)$ %(pure-read)%
• $\{\}\ r := x\ \{x = *r\}$ %(read-write)%
• $\{\neg r = s \wedge x = *r\}\ s := y\ \{x = *r\}$ %(read-write-other)%

spec DYNAMICREFERENCE = REFERENCE **then**
var $a, b, c\colon FlatCpo$
op $new\colon a \xrightarrow{c} R\ (Ref\ a)$
var $x\colon a;\ y\colon b;\ z\colon b \to c;\ r\colon Ref\ c$;
$p\colon Ref\ a \to R\ b$
• $\{\}\ r \leftarrow new\ x\ \{x = *r\}$ %(read-new)%
• $\{x = *r\}\ s \leftarrow new\ y\ \{\neg r = s \Rightarrow x = *r\}$ %(read-new-other)%
• $\{\}\ r \leftarrow new\ x;\ w \leftarrow p\ r;$
$s \leftarrow new\ (z\ w)\ \{\neg r = s\}$ %(new-distinct)%

Fig. 17. Specification of the reference and the dynamic reference monad

value. By contrast, writing to a reference does not change the values of *other* references (*read-write-other*). Note that nothing is said about the nature of references; they could e.g. be integers. The specification of *dynamic* references additionally provides an operation *new* for dynamically creating new references. Axiom *read-new* states that after initialising a reference, we can read the initial value. Moreover, creation of new references does not change the values of other references (*read-new-other*). Finally, two newly created references are distinct (*new-distinct*). Note that we do not say anything about reading from references that have not been created yet. In the discussion below, references to rules always refer to the Hoare calculus of Fig. 16.

Using this axiomatisation, we now show

$$\{\}\ r \leftarrow new\ x; s \leftarrow new\ y\ \{\neg r = s \wedge x = *r \wedge y = *s\}. \tag{1}$$

We proceed as follows. By *read-new* and rules (*Logical*) and (seq), we have

$$\{\}\ r \leftarrow new\ x; s \leftarrow new\ y\ \{y = *s\}.$$

By applying rule (seq) to *read-new* and *read-new-other*, we obtain

$$\{\}\ r \leftarrow new\ x; s \leftarrow new\ y\ \{\neg r = s \Rightarrow x = *r\}.$$

Instantiating *new-distinct* with $p = \lambda_ \bullet\ ret\ ()$ and $z = \lambda_ \bullet\ y$ and applying rule (η), we have

$$\{\}\ r \leftarrow new\ x; s \leftarrow new\ y\ \{\neg r = s\}.$$

We then obtain (1) by propositional reasoning with these three formulas.

Another example is the nondeterminism monad, shown in Fig. 18. While *fail* yields no result and hence satisfies arbitrary postconditions, *chaos* yields any result and hence nothing can be said about it. The operation $[\!]$ is nondeterministic choice (i.e. takes the union of value sets), and *sync* synchronises two nondeterministic values (i.e. takes the intersection of value sets).

spec NONDETERMINISM = FLATCPOMONAD **then**
var $a : FlatCpo$
ops $fail, chaos : N\ a;$
$_[\!]_, _sync_ : (N\ a \times N\ a) \xrightarrow{c} N\ a$
var $x : a;\ p, q : N\ a;\ \varphi, \psi : N\ Logical;$
$\chi_1, \chi_2 : a \rightarrow N\ Logical$
- $\{\}\ fail\ \{\psi\}$ %(fail)%
- $\{\varphi\}\ x \leftarrow p\ \{\chi_1 x\}\ \wedge\ \{\varphi\}\ x \leftarrow q\ \{\chi_2 x\}\ \Rightarrow$
  $\{\varphi\}\ x \leftarrow p[\!]q\ \{\chi_1 x \vee \chi_2 x\}$ %(join)%
- $\{\varphi\}\ x \leftarrow p\ \{\chi_1 x\}\ \wedge\ \{\varphi\}\ x \leftarrow q\ \{\chi_2 x\}\ \Rightarrow$
  $\{\varphi\}\ x \leftarrow p\ sync\ q\ \{\chi_1 x \wedge \chi_2 x\}$ %(sync)%

Fig. 18. The nondeterminism monad

One advantage of the looseness of the specifications introduced so far is that we now can combine the specification of references and of nondeterminism and get a specification of nondeterministic reference computations (Fig. 19).

spec NONDETERMINISTICDYNAMICREFERENCE =
DYNAMICREFERENCE **with** $R \mapsto NR$
and NONDETERMINISM **with** $N \mapsto NR$

Fig. 19. The nondeterministic dynamic reference monad

As an example, we prove the partial correctness of Dijkstra's nondeterministic version of Euclid's algorithm for computing the greatest common divisor [17] within this monad. Let *euclid* be the program sequence (over *NR Int*)

$r \leftarrow new\ x;$
$s \leftarrow new\ y;$
$while\ ret\ (\neg{*r} == {*s})$
$\quad (if\ ret\ ({*r} > {*s})\ then\ r := {*r} - {*s}\ else\ fail$
$\quad []$
$\quad if\ ret\ ({*s} > {*r})\ then\ s := {*s} - {*r}\ else\ fail)$

Assuming that we have some specification of arithmetic, including *gcd* specified to be the greatest common divisor function, we now prove

$$\{\}\ euclid\ \{{*r} = gcd(x, y)\}.$$

We proceed as follows. Using (pure-read), (comm), (copy), and propositional reasoning, we can show

$$\begin{array}{l}\{\neg r = s \wedge gcd({*r}, {*s}) = gcd(x, y) \wedge {*r} > {*s}\} \\ u \leftarrow *\ r; v \leftarrow *\ s \\ \{\neg r = s \wedge gcd({*r}, {*s}) = gcd(x, y) \wedge {*r} > {*s} \wedge u = {*r} \wedge v = {*s}\}.\end{array}$$

By congruence reasoning and (wk), we obtain

$$\begin{array}{l}\{\neg r = s \wedge gcd({*r}, {*s}) = gcd(x, y) \wedge {*r} > {*s}\} \\ u \leftarrow *\ r; v \leftarrow *\ s \\ \{\neg r = s \wedge gcd(u, v) = gcd(x, y) \wedge u > v \wedge v = {*s}\}.\end{array} \tag{2}$$

From *read-write* and *read-write-other*, we show by propositional reasoning

$$\begin{array}{l}\{\neg r = s \wedge gcd(u, v) = gcd(x, y) \wedge u > v \wedge v = {*s}\} \\ r := u - v \\ \{\neg r = s \wedge gcd(u, v) = gcd(x, y) \wedge u > v \wedge v = {*s} \wedge u - v = {*r}\}.\end{array}$$

By arithmetic reasoning and (wk), we obtain

$$\begin{array}{l}\{\neg r = s \wedge gcd(u, v) = gcd(x, y) \wedge u > v \wedge v = {*s}\} \\ r := u - v \\ \{\neg r = s \wedge gcd({*r}, {*s}) = gcd(x, y)\}.\end{array}$$

By sequencing with (2) and noting that $r := {*r} - {*s}$ is shorthand for $u \leftarrow *\ r; v \leftarrow *\ s; r := u - v$, we arrive at

$$\begin{array}{l}\{\neg r = s \wedge gcd({*r}, {*s}) = gcd(x, y) \wedge {*r} > {*s}\} \\ r := {*r} - {*s} \\ \{\neg r = s \wedge gcd({*r}, {*s}) = gcd(x, y)\}.\end{array}$$

By *fail*, we have

$$\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y) \wedge \neg *r > *s\}$$
$$fail$$
$$\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y)\}.$$

Hence by the (if) rule for pure conditions,

$$\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y)\}$$
$$if\ *r > *s\ then\ r := *r - *s\ else\ fail$$
$$\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y)\}.$$

Analogously, we have

$$\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y)\}$$
$$if\ *s > *r\ then\ s := *s - *r\ else\ fail$$
$$\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y)\}.$$

From these, we obtain by *join* and rule (wk)

$$\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y)\}$$
$$if\ ret\ (*r > *s)\ then\ r := *r - *s\ else\ fail$$
$$[\!]\ if\ ret\ (*s > *r)\ then\ s := *s - *r\ else\ fail$$
$$\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y)\}.$$

Applying the standard (while) rule and rule (wk) leads to

$$\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y)\}$$
$$while\ ret\ (\neg *r == *s)$$
$$(\ \ if\ ret\ (*r > *s)\ then\ r := *r - *s\ else\ fail$$
$$[\!]\ if\ ret\ (*s > *r)\ then\ s := *s - *r\ else\ fail)$$
$$\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y) \wedge *r == *s\}.$$

Using the arithmetic fact that $gcd(z, z) = z$, we obtain by (wk)

$$\begin{aligned}
&\{\neg r = s \wedge gcd(*r, *s) = gcd(x, y)\} \\
&while\ \neg *r == *s \\
&(\ \ if\ ret\ (*r > *s)\ then\ r := *r - *s\ else\ fail \\
&[\!]\ if\ ret\ (*s > *r)\ then\ s := *s - *r\ else\ fail) \\
&\{*r = gcd(x, y)\}.
\end{aligned} \tag{3}$$

From (1) above, we obtain by congruence reasoning and rule (wk)

$$\{\}\ r \leftarrow new\ x; s \leftarrow new\ y\ \{\neg r = s \wedge gcd(*r, *s) = gcd(x, y)\}, \tag{4}$$

and the result now follows by sequencing (3) and (4).

12 Conclusions

We have presented the design of HasCasl, a wide-spectrum language serving the integrated specification and development of software as well as mathematical modelling on a wider scale. Novel features of HasCasl include the semantic treatment of type class polymorphism by means of an extension semantics and support for inductive datatypes and recursion that does not rely on unique choice. We have moreover laid out the technical aspects of the syntax of the type class mechanism and its interaction with higher order subtyping in some detail.

We have illustrated the expressive strength of HasCasl by means of the development of a Hoare logic for monad-encapsulated generic side effects as used in modern functional-imperative programming. The latter is a contribution in its own right, as it offers modularised reasoning support for monad-based imperative programs, where generic rules are cleanly separated from axiomatisations of specific notions of side-effect (i.e. monads). A stronger generic computational logic of this nature, namely a monad-based dynamic logic, has been presented in [75,52]; this extension, however, relies on stronger assumptions on the underlying monad. The use of HasCasl outside the realm of software specification as such has been illustrated in [81], where composition tables of region connection calculi are verified in a logically heterogeneous setting in which HasCasl serves the definition of higher-order concepts such as the real numbers.

HasCasl is a central node in the logic graph of the Bremen heterogeneous tool set. As such, it is provided with extensible reasoning support, presently via a translation into Isabelle/HOL, and further connections to other logics in the graph, e.g. a translation of executable specifications to Haskell. These tools are being developed further; in particular, the technical handling of Isabelle proofs on translated HasCasl specifications and the development of suitable dedicated tactics is the subject of ongoing work. Experimental work on the verification of the Haskell prelude against a HasCasl specification is in progress [9].

An open issue in the language design of HasCasl itself is the specification of nested polymorphism as supported by the Glasgow extensions of Haskell [59], i.e. to find a workaround for the fact that higher order logic is inconsistent with System F [15]. An initial step in this direction would be the support for existential types, which provide a clean way of encapsulating representations of abstract datatypes [37]. Concerning the HasCasl development methodology (and indeed any methodology that works with standard logics on non-continuous functions to specify higher order programs), an open problem is support for developments that start with an abstract algebraic specification

and only later refine this to a design specification working with cpo's and continuous functions. Currently, this works smoothly only for first-order functions, whereas for many higher-order functions, one has to work with continuous function spaces from the outset in order to avoid possible dead ends (in fact, this problem is not even entirely solved by working with continuous functions, as any actual implementation will interpret higher-order types as types of computable functions).

Some of the general limitations of the algebraic methodology of program specification listed in [69] also apply to HasCasl: firstly, the relation to informal requirements is not addressed in the current work — we consider this to be an important but separate issue. Secondly, real programming languages often have subtle complex features that are ignored in the specification world; one example that concerns HasCasl is Haskell's lazy pattern matching (which, however, could be integrated without too much effort using the work in [27]). Finally, while the monad-based approach does support object-oriented and concurrent programming as shown in work on Haskell, experiments are needed that establish the feasibility of verification involving these features in HasCasl.

Acknowledgements

The authors wish to thank Mihai Codescu, Kathrin Hoffmann, Bernd Krieg-Brückner, Christoph Lüth, Christian Maeder, and Don Sannella for useful comments and discussions, and the referees for their valuable suggestions. Erwin R. Catesbeiana contributed his opinion on matters of consistency.

References

[1] A. Abel and R. Mattes. Fixed points of type constructors and primitive recursion. In J. Marcinkowski and A. Tarlecki, eds., *Computer Science Logic, CSL 2004*, vol. 3210 of *Lect. Notes Comput. Sci.*, pp. 190–204. Springer, 2004.

[2] J. Adámek, H. Herrlich, and G. E. Strecker. *Abstract and Concrete Categories.* Wiley Interscience, 1990.

[3] E. Astesiano and M. Cerioli. Free objects and equational deduction for partial conditional specifications. *Theoret. Comput. Sci.*, 152:91–138, 1995.

[4] J. Bates and R. Constable. Proofs as programs. *ACM Trans. Prog. Lang. Systems*, 7:113–136, 1985.

[5] S. Berghofer and M. Wenzel. Inductive datatypes in HOL - lessons learned in formal-logic engineering. In Y. Bertot, G. Dowek, A. Hirschowitz, C. Paulin, and L. Théry, eds., *Theorem Proving in Higher Order Logics, TPHOLs 1999*, vol. 1690 of *Lect. Notes Comput. Sci.*, pp. 19–36. Springer, 1999.

[6] M. Bidoit and P. D. Mosses. CASL *User Manual*, vol. 2900 of *Lect. Notes Comput. Sci.* Springer, 2004.

[7] L. Birkedal and R. E. Møgelberg. Categorical models for Abadi and Plotkin's logic for parametricity. *Math. Struct. Comput. Sci.*, 15, 2005.

[8] M. Broy, C. Facchi, R. Grosu, R. Hettler, H. Hussmann, D. Nazareth, F. Regensburger, and K. Stølen. The requirement and design specification language Spectrum, an informal introduction, version 1.0. Technical report, Institut für Informatik, Technische Universität München, Mar. 1993.

[9] G. M. Cabral. Developing a HASCASL library for the Haskell prelude. Master's thesis, Institute of Computing, University of Campinas. In preparation.

[10] L. Cardelli. Notes on $F^{\omega}_{<:}$. Unpublished notes, 1990.

[11] M. Cerioli and J. Meseguer. May I borrow your logic? (transporting logical structures along maps). *Theoret. Comput. Sci.*, 173:311–347, 1997.

[12] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. Talcott. *All About Maude - A High-Performance Logical Framework. How to Specify, Program, and Verify Systems in Rewriting Logic*, vol. 4350 of *Lect. Notes Comput. Sci.* Springer, 2007.

[13] CoFI. The Common Framework Initiative for algebraic specification and development. Electronic Archives under www.cofi.info.

[14] The Coq Development Team. The Coq Proof Assistant – Reference Manual, v8.1. INRIA, 2006, Available under http://coq.inria.fr.

[15] T. Coquand. An analysis of Girard's paradox. In *Logic in Computer Science, LICS 1986*, pp. 227–236. IEEE, 1986.

[16] R. Diaconescu and K. Futatsugi. *CafeOBJ Report.* AMAST series. World Scientific, Singapore, 1998.

[17] E. W. Dijkstra. *A Discipline of Programming.* Prentice Hall, 1976.

[18] J.-C. Filliâtre. Proof of imperative programs in type theory. In T. Altenkirch, W. Naraschewski, and B. Reus, eds., *Types for Proofs and Programs, TYPES 1998*, vol. 1657 of *Lect. Notes Comput. Sci.*, pp. 78–92. Springer, 1999.

[19] C. Führmann. Varieties of effects. In M. Nielsen and U. Engberg, eds., *Foundations of Software Science And Computation Structures, FOSSACS 2002*, vol. 2303 of *Lect. Notes Comput. Sci.*, pp. 144–158. Springer, 2002.

[20] M. J. Gabbay and A. M. Pitts. A new approach to abstract syntax with variable binding. *Formal Asp. Comput.*, 13:341–363, 2002.

[21] C. George, P. Haff, K. Havelund, A. E. Haxthausen, R. Milne, C. B. Nielson, S. Prehn, and K. R. Wagner. *The Raise Specification Language*. Prentice Hall, New York, 1992.

[22] J. Goguen and R. Burstall. Institutions: Abstract model theory for specification and programming. *J. ACM*, 39:95–146, 1992.

[23] J. A. Goguen and G. Rosu. Institution morphisms. *Formal Asp. Comput.*, 13:274–307, 2002.

[24] S. Goncharov, L. Schröder, and T. Mossakowski. Completeness of global evaluation logic. In R. Kralovic and P. Urzyczyn, eds., *Mathematical Foundations of Computer Science, MFCS 2006*, vol. 4162 of *Lect. Notes Comput. Sci.*, pp. 447–458. Springer, 2006.

[25] J. V. Guttag, J. J. Horning, S. J. Garland, K. D. Jones, A. Modet, and J. M. Wing. *Larch: Languages and Tools for Formal Specification*. Springer, 1993.

[26] T. Hallgren. Haskell tools from the programatica project. In J. Jeuring, ed., *Haskell Workshop, HASKELL 2003*, pp. 103–106. ACM Press, 2003.

[27] W. L. Harrison and R. B. Kieburtz. The logic of demand in Haskell. *J. Funct. Programming*, 15:837–891, 2005.

[28] A. Haxthausen. Order-sorted algebraic specifications with higher-order functions. *Theoret. Comput. Sci.*, 183:157–185, 1997.

[29] L. Henkin. The completeness of the first-order functional calculus. *J. Symbolic Logic*, 14:159–166, 1949.

[30] H. Herrlich, E. Lowen-Colebunders, and F. Schwarz. Improving Top: PrTop and PsTop. In H. Herrlich and H.-E. Porst, eds., *Category theory at work*, pp. 21–34. Heldermann, Berlin, 1991.

[31] D. Hutter, H. Mantel, G. Rock, W. Stephan, A. Wolpers, M. Balser, W. Reif, G. Schellhorn, and K. Stenzel. VSE: Controlling the complexity in formal software developments. In D. Hutter, W. Stephan, P. Traverso, and M. Ullmann, eds., *Current Trends in Applied Formal Methods, FM-Trends 1998*, vol. 1641 of *Lect. Notes Comput. Sci.* Springer, 1999.

[32] C. B. Jones. *Systematic Software Development Using VDM.* Prentice Hall, 1990.

[33] S. Kahrs and D. Sannella. Reflections on the design of a specification language. In M. Nivat, ed., *Fundamental Approaches to Software Engineering, FOSSACS 1998*, vol. 1382 of *Lect. Notes Comput. Sci.*, pp. 154–170. Springer, 1998.

[34] S. Kahrs, D. Sannella, and A. Tarlecki. The definition of extended ML: A gentle introduction. *Theoret. Comput. Sci.*, 173:445–484, 1997.

[35] R. B. Kieburtz. P-logic: property verification for Haskell programs, 2002. Unpublished manuscript.

[36] J. Lambek and P. J. Scott. *Introduction to Higher Order Categorical Logic.* Cambridge University Press, 1986.

[37] K. Läufer. Type classes with existential types. *J. Functional Programming*, 6:485–517, 1996.

[38] S. Liang, P. Hudak, and M. Jones. Monad transformers and modular interpreters. In *Principles of Programming Languages, POPL 95*, pp. 333–343. ACM Press, 1995.

[39] S. Mac Lane. *Categories for the Working Mathematician.* Springer, 1997.

[40] P. Martin-Löf. Constructive mathematics and computer programming. In *Logic, Methodology, and Philosophy of Science, LMPS 1979*, pp. 153–175. North-Holland, 1982.

[41] J. Meseguer. General logics. In *Logic Colloquium, LC 1987*, pp. 275–329. North Holland, 1989.

[42] G. Mints. *A Short Introduction to Intuitionistic Logic.* Kluwer, 2000.

[43] J. C. Mitchell and P. J. Scott. Typed lambda models and cartesian closed categories. In J. Gray and A. Scedrov, eds., *Categories in Computer Science and Logic*, vol. 92 of *Contemp. Math.*, pp. 301–316. Amer. Math. Soc., 1989.

[44] E. Moggi. Categories of partial morphisms and the λ_p-calculus. In D. H. Pitt, S. Abramsky, A. Poigné, and D. E. Rydeheard, eds., *Category Theory and Computer Programming*, vol. 240 of *Lect. Notes Comput. Sci.*, pp. 242–251. Springer, 1986.

[45] E. Moggi. *The Partial Lambda Calculus.* PhD thesis, University of Edinburgh, 1988.

[46] E. Moggi. Notions of computation and monads. *Inform. and Comput.*, 93:55–92, 1991.

[47] E. Moggi. A semantics for evaluation logic. *Fund. Inform.*, 22:117–152, 1995.

[48] T. Mossakowski. Relating CASL with other specification languages: the institution level. *Theoret. Comput. Sci.*, 286:367–475, 2002.

[49] T. Mossakowski. Heterogeneous specification and the heterogeneous tool set. Habilitation thesis, Universität Bremen, 2005.

[50] T. Mossakowski, C. Maeder, and K. Lüttich. The Heterogeneous Tool Set. In O. Grumberg and M. Huth, eds., *Tools and Algorithms for the Construction and Analysis of Systems, TACAS 07*, vol. 4424 of *Lect. Notes Comput. Sci.*, pp. 519–522. Springer, 2007.

[51] T. Mossakowski, D. Sannella, and A. Tarlecki. A simple refinement language for casl. In J. L. Fiadeiro, ed., *Recent Developments in Algebraic Development Techniques, 17th International Workshop, WADT 2004*, vol. 3423 of *Lect. Notes Comput. Sci.*, pp. 162–185. Springer, 2005.

[52] T. Mossakowski, L. Schröder, and S. Goncharov. A generic complete dynamic logic for reasoning about purity and effects. In J. Fiadeiro and P. Inverardi, eds., *Fundamental Approaches to Software Engineering, FASE 2008*, vol. 4961 of *Lect. Notes Comput. Sci.*, pp. 199–214. Springer, 2008.

[53] P. D. Mosses, ed. CASL *Reference Manual*, vol. 2960 of *Lect. Notes Comput. Sci.* Springer, 2004.

[54] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, vol. 2283 of *Lect. Notes Comput. Sci.* Springer, 2002.

[55] J. Nordlander. Polymorphic subtyping in O'Haskell. *Sci. Comput. Programming*, 43(2-3):93–127, 2002.

[56] S. Owre, N. Shankar, J. M. Rushby, and D. W. J. Stringer-Calvert. *PVS Language Reference, Version 2.4.* SRI International, Menlo Park, 2001.

[57] L. C. Paulson. Mechanizing coinduction and corecursion in higher-order logic. *J. Log. Comput*, 7:175–204, 1997.

[58] S. Peyton Jones, ed. *Haskell 98 Language and Libraries — The Revised Report.* Cambridge, 2003. also: *J. Funct. Programming* 13, 2003.

[59] S. Peyton Jones, D. Vytiniotis, S. Weirich, and M. Shields. Practical type inference for arbitrary-rank types. *J. Funct. Programming*, 17:1–82, 2007.

[60] W. Phoa. An introduction to fibrations, topos theory, the effective topos and modest sets. Research report ECS-LFCS-92-208, Lab. for Foundations of Computer Science, University of Edinburgh, 1992.

[61] B. Pierce. *Types and Programming Languages.* MIT Press, 2002.

[62] A. Pitts. Evaluation logic. In G. Birtwhistle, ed., *Higher Order Workshop IV*, Workshops in Computing, pp. 162–189. Springer, 1991.

[63] A. M. Pitts. Polymorphism is set-theoretic, constructively. In *Category Theory and Computer Science*, vol. 283 of *Lect. Notes Comput. Sci.*, pp. 12–39. Springer, 1987.

[64] F. Regensburger. HOLCF: Higher order logic of computable functions. In E. T. Schubert, P. J. Windley, and J. Alves-Foss, eds., *Theorem Proving in Higher Order Logics, TPHOLS 1995*, vol. 971 of *Lect. Notes Comput. Sci.*, pp. 293–307, 1995.

[65] G. Rosolini. *Continuity and effectiveness in topoi.* PhD thesis, University of Oxford, 1986.

[66] G. Rosolini and T. Streicher. Comparing models of higher type computation. In *Realizability Semantics and Applications*, vol. 23 of *Electron. Notes Theoret. Comput. Sci.*, 1999.

[67] D. Sannella and A. Tarlecki. *Foundations of Algebraic Specifications and Formal Program Development.* In preparation. Preliminary version available under http://homepages.inf.ed.ac.uk/dts/book.

[68] D. Sannella and A. Tarlecki. Toward formal development of programs from algebraic specifications: implementations revisited. *Acta Inf.*, 25:233–281, 1988.

[69] D. Sannella and A. Tarlecki. Algebraic methods for specification and formal development of programs. *ACM Comput. Surveys*, 31, 1999.

[70] L. Schröder. The logic of the partial λ-calculus with equality. In J. Marcinkowski and A. Tarlecki, eds., *Computer Science Logic, CSL 2004*, vol. 3210 of *Lect. Notes Comput. Sci.*, pp. 385–399. Springer, 2004.

[71] L. Schröder. The HasCasl prologue - categorical syntax and semantics of the partial λ-calculus. *Theoret. Comput. Sci.*, 353:1–25, 2006.

[72] L. Schröder. Bootstrapping types and cotypes in HasCasl. In T. Mossakowski and U. Montanari, eds., *Algebra and Coalgebra in Computer Science, CALCO 2007*, vol. 4624 of *Lect. Notes Comput. Sci.*, pp. 447–462. Springer, 2007. Full version to appear in *Log. Methods Comput. Sci.*

[73] L. Schröder and T. Mossakowski. HasCasl: Towards integrated specification and development of Haskell programs. In H. Kirchner and C. Ringeissen, eds., *Algebraic Methodology and Software Technology, AMAST 2002*, vol. 2422 of *Lect. Notes Comput. Sci.*, pp. 99–116. Springer, 2002.

[74] L. Schröder and T. Mossakowski. Monad-independent Hoare logic in HasCasl. In M. Pezzè, ed., *Fundamental Approaches to Software Engineering, FASE 2003*, vol. 2621 of *Lect. Notes Comput. Sci.*, pp. 261–277. Springer, 2003.

[75] L. Schröder and T. Mossakowski. Monad-independent dynamic logic in HasCasl. *J. Logic Comput.*, 14:571–619, 2004.

[76] L. Schröder, T. Mossakowski, and C. Lüth. Type class polymorphism in an institutional framework. In J. Fiadeiro, ed., *Recent Developments in Algebraic Development Techniques, 17th International Workshop, WADT 04*, vol. 3423 of *Lect. Notes Comput. Sci.*, pp. 234–248. Springer, 2004.

[77] L. Schröder, T. Mossakowski, A. Tarlecki, P. Hoffman, and B. Klin. Semantics of architectural specifications in Casl. In H. Hußmann, ed., *Fundamental Approaches to Software Engineering, FASE 2001*, vol. 2029 of *Lect. Notes Comput. Sci.*, pp. 253–268. Springer, 2001.

[78] M. Spivey. *The Z Notation: A Reference Manual.* Prentice Hall, 1992. 2nd edition.

[79] H. Thielecke. *Categorical Structure of Continuation Passing Style.* PhD thesis, University of Edinburgh, 1997.

[80] P. Wadler. How to declare an imperative. *ACM Computing Surveys*, 29:240–263, 1997.

[81] S. Wölfl, T. Mossakowski, and L. Schröder. Qualitative constraint calculi: Heterogeneous verification of composition tables. In D. Wilson and G. Sutcliffe, eds., *20th International FLAIRS Conference*, pp. 665–670. AAAI Press, 2007.

[82] O. Wyler. *Lecture notes on topoi and quasitopoi.* World Scientific, 1991.

Appendix

A Subkinding Rules

For convenience, the full set of subkinding rules as assembled in Sect.s 3 and 4 is shown in Fig. A.1. Recall that $\mathcal{V}$ denotes the set $\{\pm, +, -, \cdot\}$ of variance annotations, ordered by taking $\pm$ and $\cdot$ to be the smallest and the greatest element, respectively, and $+$ and $-$ to be incomparable.

$$\frac{Cl \leq_C Kd \text{ in } \Sigma}{Cl \leq_K Kd} \qquad \frac{Kd_1 \leq_K Kd_2 \quad Kd_3 \leq_K Kd_4}{\mu Kd_2 \to Kd_3 \leq_K \mu Kd_1 \to Kd_4} \ (\mu \in \mathcal{V})$$

$$\frac{}{\mu Kd_1 \to Kd_2 \leq_K \nu Kd_1 \to Kd_2} \ (\mu, \nu \in \mathcal{V}, \mu \leq \nu),$$

$$\frac{}{Kd \leq_K Kd} \qquad \frac{Kd_1 \leq_K Kd_2 \quad Kd_2 \leq_K Kd_3}{Kd_1 \leq_K Kd_3}$$

Fig. A.1. Subkinding rules

B Kinding Rules

The full set of kinding rules for pseudotypes as assembled in Sect. 3 and 4 is shown in Fig. B.1. Recall that Θ^{-1} and Θ^0 denote the context Θ with all outer variances reversed or removed, respectively.

$$\frac{F : Kd_1 \text{ in } \Sigma \quad Kd_1 \leq_K Kd_2}{\Theta \rhd F : Kd_2} \qquad \frac{a : \mu Kd_1 \text{ in } \Theta, \mu \in \{+, \cdot\} \quad Kd_1 \leq_K Kd_2}{\Theta \rhd a : Kd_2}$$

$$\frac{\Theta^0 \rhd t : Kd_1 \quad \Theta \rhd s : Kd_1 \to Kd_2}{\Theta \rhd s\ t : Kd_2} \qquad \frac{\Theta \rhd t : Kd_1 \quad \Theta \rhd s : +Kd_1 \to Kd_2}{\Theta \rhd s\ t : Kd_2} \qquad \frac{\Theta^{-1} \rhd t : Kd_1 \quad \Theta \rhd s : -Kd_1 \to Kd_2}{\Theta \rhd s\ t : Kd_2}$$

$$\frac{\Theta, a : \mu Kd_1 \rhd t : Kd_2 \quad Kd_3 \leq_K Kd_1}{\Theta \rhd \lambda\, a : Kd_1 \bullet t : \nu Kd_3 \to Kd_2} \ (\mu \leq \nu \text{ in } \mathcal{V})$$

Fig. B.1. Kinding rules for type constructors

C Syntax-directed Subkinding Rules

For implementation purposes, we give a syntax-directed version of the subkinding rules (Appendix A). The point is to eliminate the transitivity and reflexivity rules in the spirit of 'algorithmic subtyping' [61]. The syntax-directed rules are given in Fig. C.1.

$$\textbf{(cl-refl)}\ \frac{}{Cl \leq_K Cl} \qquad \textbf{(cl)}\ \frac{\begin{array}{c} Cl \leq_C Kd_1 \\ Kd_1 \leq_K Kd_2 \end{array}}{Cl \leq_K Kd_2}$$

$$(\rightarrow)\ \frac{Kd_1 \leq_K Kd_2 \quad Kd_3 \leq_K Kd_4}{\mu Kd_2 \rightarrow Kd_3 \leq_K \nu Kd_1 \rightarrow Kd_4}\ (\mu, \nu \in \mathcal{V}, \mu \leq \nu)$$

Fig. C.1. Syntax-directed subkinding rules

Note that rule (cl) is indeed algorithmic since there are only finitely many declarations $Cl \leq_C Kd_1$. In rule $(\rightarrow)$, $\mathcal{V}$ is the set of variance annotations, ordered as described in Sect. 4.

Proposition 75 *The rules of Fig. C.1 derive the same subkinding judgements as the rules given in Fig. A.1 above.*

PROOF. (Sketch) It is clear that the rules of Fig. C.1 are derivable from the previous rules and subsume all of the previous rules except reflexivity and transitivity. By induction over the kind structure, it is easy to show that $Kd \leq_K Kd$ is derivable by the rules of Fig. C.1 for all kinds Kd. Finally, the fact that the relation $\leq_K$ generated by the rules of Fig. C.1 and the reflexivity rule is transitive is shown by induction over the combined lengths of the derivations of $Kd_1 \leq_K Kd_2$ and $Kd_2 \leq_K Kd_3$; this involves a case distinction over which rules were applied in the last step in either case.

D Syntax-directed Subtyping Rules

Similarly as for the subkinding system, one can give a syntax-directed set of rules, shown in Fig. D.1, for the subtype relation which is equivalent to the rules presented in Sect. 4. The proof of equivalence is analogous to the one sketched for Prop. 75. The introduction rules for variables and type constructors are, like the rule (cl) of Fig. C.1, algorithmic because there are only finitely many declarations $F \leq t$ and $a \leq t$ in Σ and Λ, respectively. As in Fig. 7, $\sqsubseteq$ ranges over $\{\leq, \leq_*\}$.

$$\frac{a \text{ in } \Theta}{\Theta; \Lambda \rhd a \sqsubseteq a} \qquad \frac{}{\Theta; \Lambda \rhd F \sqsubseteq F} \qquad \frac{\begin{array}{c} a \leq t \text{ in } \Lambda \\ \Theta; \Lambda \rhd t \sqsubseteq s \end{array}}{\Theta; \Lambda \rhd a \sqsubseteq s} \qquad \frac{\begin{array}{c} F \leq t \text{ in } \Sigma \\ \Theta; \Lambda \rhd t \sqsubseteq s \end{array}}{\Theta; \Lambda \rhd F \sqsubseteq s}$$

$$\frac{\begin{array}{c} \Theta \rhd t_1, t_2 : +Kd_1 \to Kd_2 \\ \Theta; \Lambda \rhd s_1 \sqsubseteq s_2 \\ \Theta; \Lambda \rhd t_1 \sqsubseteq t_2 \end{array}}{\Theta; \Lambda \rhd t_1\ s_1 \sqsubseteq t_2\ s_2} \qquad \frac{\begin{array}{c} \Theta \rhd t_1, t_2 : -Kd_1 \to Kd_2 \\ \Theta; \Lambda \rhd s_2 \leq_* s_1 \\ \Theta; \Lambda \rhd t_1 \leq_* t_2 \end{array}}{\Theta; \Lambda \rhd t_1\ s_1 \leq_* t_2\ s_2}$$

$$\frac{\Theta, a : Kd; \Lambda \rhd t \sqsubseteq s}{\Theta; \Lambda \rhd \lambda\, a : \mu Kd \bullet t \sqsubseteq \lambda\, a : \mu Kd \bullet s}\ (\mu \in \mathcal{V})$$

Fig. D.1. Syntax-directed subtyping rules for pseudotypes