

Executable Component-Based Semantics

L. Thomas van Binsbergen^{a,*}, Peter D. Mosses^{b,c}, Neil Sculthorpe^d

^a*Department of Computer Science, Royal Holloway, University of London,
TW20 0EX, Egham, United Kingdom*

^b*Department of Computer Science, Swansea University,
SA2 8PP, Swansea, United Kingdom*

^c*EEMCS, Programming Languages, Delft University of Technology,
P.O. Box 5031, 2600 GA Delft, The Netherlands*

^d*Department of Computing and Technology, Nottingham Trent University,
NG11 8NS, Nottingham, United Kingdom*

Abstract

The potential benefits of formal semantics are well known. However, a substantial amount of work is required to produce a complete and accurate formal semantics for a major language; and when the language evolves, large-scale revision of the semantics may be needed to reflect the changes. The investment of effort needed to produce an initial definition, and subsequently to revise it, has discouraged language developers from using formal semantics. Consequently, many major programming languages (and most domain-specific languages) do not yet have formal semantic definitions.

To improve the practicality of formal semantic definitions, the PLANCOMPS project has developed a component-based approach. In this approach, the semantics of a language is defined by translating its constructs (compositionally) to combinations of so-called fundamental constructs, or ‘funcons’. Each funcon is defined using a modular variant of Structural Operational Semantics, and forms a language-independent component that can be reused in definitions of different languages. A substantial library of funcons has been developed and tested in several case studies. Crucially, the definition of each funcon is fixed, and does not need changing when new funcons are added to the library.

For specifying component-based semantics, we have designed and implemented a meta-language called CBS. It includes specification of abstract syntax, of its translation to funcons, and of the funcons themselves. Development of CBS specifications is supported by an integrated development environment. The accuracy of a language definition can be tested by executing the specified translation on programs written in the defined language, and then executing the resulting funcon terms using an interpreter generated from the CBS definitions of the funcons. This paper gives an introduction to CBS, illustrates its use, and presents the various tools involved in our implementation of CBS.

Keywords: programming languages, formal semantics, reuse, components, tool support

*Corresponding author

Email addresses: ltvanbinsbergen@acm.org (L. Thomas van Binsbergen), p.d.mosses@swansea.ac.uk (Peter D. Mosses), neil.sculthorpe@ntu.ac.uk (Neil Sculthorpe)

1. Introduction

New programming languages and domain-specific languages are continually being introduced, as are new versions of existing languages. Each language needs to be carefully specified, to determine the syntax and semantics of its programs. Context-free aspects of syntax are usually specified, precisely and succinctly, using formal grammars; in contrast, semantics (including static checks and disambiguation) is generally specified only informally, without use of precise notation. Informal specifications are often incomplete or inconsistent, and open to misinterpretation; formal specifications can avoid such issues. Moreover, completely formal definitions of programming languages may be used to generate prototype implementations, and as a basis for proving properties of languages and of individual programs.

Although there is broad agreement as to the benefits of formality in language definitions, and although there are a few examples of successful individual projects (notably the definition of STANDARD ML [1]), there is generally little inclination on the part of programming language developers themselves to produce a complete formal semantics definition. It appears that this is at least partly due to the effort required when scaling up to larger languages, and when updating a formal semantics to reflect language evolution (see, for instance, comments by the HASKELL designers [2, Page 9]).

New languages typically include a large number of constructs from previous languages, presenting a major opportunity for reuse of specification components. However, in the absence of a suitable collection of reusable components, each language would have to be specified from scratch—a huge effort.

1.1. Component-Based Semantics

To improve the practicality of formal semantic definitions of larger languages, the PLANCOMPS project¹ has developed a component-based approach to semantics. In this approach, a reusable component of language definitions corresponds to a *fundamental programming construct*: a so-called ‘*funcon*’, which has a fixed operational interpretation. The formal semantics of each funcon is defined independently, using I-MSOS [3], a variant of Modular Structural Operational Semantics [4]. For example, the funcon **if-true-else** is defined as follows:

Funcon **if-true-else**($_ : \mathbf{booleans}$, $_ : \Rightarrow T$, $_ : \Rightarrow T$) : $\Rightarrow T$

Rule **if-true-else**(**true**, X , $_$) $\rightsquigarrow X$

Rule **if-true-else**(**false**, $_$, Y) $\rightsquigarrow Y$

The collection of funcons is open-ended; crucially, adding new funcons does not require changes to the definition or use of previous funcons.

A component-based semantics of a programming language is defined by translating its constructs to funcons. For example, a non-strict conjunction with Boolean-valued operands could be translated as follows:

Rule $\mathit{rval} \llbracket \mathit{Exp}_1 \text{ ' \&\&' } \mathit{Exp}_2 \rrbracket = \mathbf{if-true-else}(\mathit{rval} \llbracket \mathit{Exp}_1 \rrbracket, \mathit{rval} \llbracket \mathit{Exp}_2 \rrbracket, \mathbf{false})$

Many funcons can be widely reused in the definitions of different languages. An initial case study [5] gave a semantics for CAML LIGHT [6] based on a preliminary collection of funcons. The reusability of funcons has subsequently been validated by several further case studies, some of which are available online [7], including an update of the CAML LIGHT specification to OCAML LIGHT. These case studies demonstrate how translation to funcons scales up for medium-sized languages. A further case study (C#) has been initiated to test how well CBS can cope with a major programming language.

Analogous practices are widely adopted in software engineering: developers rely on reusable components in the form of packages. However, applications generally rely on the details of particular versions of packages, and problems can arise when new versions of packages are installed, requiring changes to applications that use them. In contrast, each individual funcon definition is fixed, and language definitions require changes only when the language itself evolves. For example, the above translation of the expression “ $\mathit{Exp}_1 \text{ ' \&\&' } \mathit{Exp}_2$ ”

¹<http://plancomps.org>

would need changing if conjunction were to become strict, or if its arguments were no longer required to be Booleans, since the definition of the **if-true-else** funcon cannot change.

This article presents CBS, a unified meta-language for component-based semantics of programming languages. CBS includes abstract syntax grammars (essentially BNF with regular expressions), the signatures and equations for functions translating language constructs to funcons, and the signatures and rules for defining funcons.

1.2. Executability of CBS

One of the most important properties of a language definition is whether the defined language corresponds to the intentions of the language developers. In practice, the only way of checking this property is to take a collection of test programs, and compare the intended behaviour of each program with that determined by the language definition. Provided that the collection of test programs is sufficiently representative, and that their behaviour involves all parts of the language definition, such empirical validation can detect most discrepancies between the intended and the defined behaviour, and establish a high degree of confidence in the accuracy of the language definition.

In general, however, the program behaviour determined by a language definition is not immediately apparent. One approach is to carry out formal proofs of the relationship between programs and their behaviours, based on facts derived from the language definition. A less demanding approach is to derive an *implementation* of the language from its definition, and observe the behaviour of programs when run. When the derived implementation can be generated automatically from the language definition, the latter is called *executable*.

We have implemented software tools to make component-based semantic definitions in CBS executable. CBS definitions are developed and maintained within an IDE (integrated development environment) that offers syntax highlighting, clickable references between uses of funcons and their definitions, as well as access to the language-specific tools generated from CBS definitions. From a CBS definition, a user can generate a parser, a translator for translating abstract syntax trees to funcon terms, and an interpreter for executing funcon terms. In combination, these generated artifacts form an interpreter for the defined language.

1.3. Novel Features of CBS

This article presents details of the CBS meta-language, and shows how it is used to formally define languages as well as an extensible library of reusable programming constructs. CBS is novel in its emphasis on reuse: language definitions are based on an extensible collection of predefined, fixed funcons, which are themselves defined in CBS. The library of reusable funcons provided with CBS is moderated, and a new funcon is added to the library only after its definition has been validated through analysis and experimentation. Language definitions reuse funcons simply by referring to their names; there are no explicit imports of modules or packages, nor any mention of the files or folders in which definitions are stored. This is possible because a particular funcon name always refers to the same definition within and across language definitions. Implemented as a Spoofox project [8], CBS provides name resolution and (clickable) references between occurrences of funcons and their definitions. The moderated collection of reusable funcons is still under development, and has been pre-released [7] to encourage discussion and review.

If necessary, language-specific funcons can be defined as part of a language definition. Such funcon definitions are modified freely, and are not intended for reuse. However, a generalisation of a language-specific funcon may reveal itself as a useful reusable component.

Funcons are highly modular: each funcon is defined independently, and its definition can be understood in isolation. Funcons inherit their modularity from the I-MSOS inference rules that define them, which involve the implicit propagation of several classes of auxiliary semantic entities. The I-MSOS rules used in CBS are based on a rich variety of such classes. Since funcons are defined with small-step style rules, it is relatively simple to give the semantics of constructs that involve interleaving, or abnormal control flow such as delimited continuations, exceptions and return statements.

1.4. Contributions and Overview

This paper contributes by:

- Introducing CBS, a novel specification language with an extensible collection of fixed, highly reusable funcons for modular, component-based language definitions.
- Demonstrating how an existing language is defined using CBS, and how to extend a language definition in a modular fashion when the language evolves.
- Introducing two new classes of auxiliary semantic entities, along with illustrative examples of their use.
- Explaining how funcon definitions can be implemented as micro-interpreters, which are as modular as the definitions from which they are generated.

As a running example, we use CBS to formally define SIMPLE, a small C-like imperative programming language [9]. The SIMPLE definition involves funcons that are likely to be reused in the definitions of most programming languages.

The rest of this article is organised as follows. In Sect. 2 we discuss the background for CBS, and related frameworks. In Sect. 3 we show how to use CBS to define abstract syntax and translation of abstract syntax trees to funcons. In Sect. 4 we show the definitions of some of the funcons that occur in the translations of the SIMPLE constructs, and explain the different classes of semantic entities used in these definitions. In Sect. 5, the SIMPLE language is extended with a delimited control operator to demonstrate the evolution of CBS language definitions as well as the novel bidirectional signals. In Sect. 6 we give an overview of the Haskell framework for implementing funcon definitions. In Sect. 7 we present an IDE for developing and validating CBS language definitions.

Appendix A provides a complete CBS definition of SIMPLE (omitting concurrency constructs) and Appendix B lists the signatures of all the funcons used in the definition. The CBS definitions of these funcons are available online [7].

2. Background and Related Work

In this section we discuss the background for CBS, and its relation to other current frameworks for formal semantics that have a high degree of modularity as well as executable definitions.

2.1. Theoretical Foundations

As mentioned in the Introduction, CBS includes meta-notation for specifying abstract syntax, translation functions (from abstract syntax to funcons), and funcons. The details of CBS are explained and illustrated in subsequent sections; here, we give an overview of the foundations of the framework.

Grammars for specifying abstract syntax in CBS are context-free. They are expressed in a variant of Extended BNF, using standard regular expression notation for sequences, alternatives, and optional and iterated phrases. Abstract syntax trees (ASTs) correspond to finite terms constructed by operations corresponding to the alternatives for each sort of phrase. (CBS also supports disambiguation rules for parsing based on abstract syntax grammars, but that is not significant for its foundations.)

CBS allows multiple translation functions for each sort of AST, defined inductively by mutual recursion. The foundations of inductive definitions of families of functions are well known.

The rules used to specify behaviour in CBS resemble rules in conventional Structural Operational Semantics [10]. However, their interpretation is based on a combination of value-computation systems [11] and Implicitly-Modular SOS (I-MSOS) [3]:

- Value-computation systems provide a context-insensitive rewriting relation as well as a transition relation, and allow transitions to be labelled by multiple entities as in Modular SOS (MSOS) [4]. The format of the rules ensures that (strong) bisimulation is a congruence, and that it is preserved by

disjoint extension. The distinction between values and computations is crucial for allowing nested patterns in source terms of transition rules (e.g., the constants in the above rules for **‘if-true-else’**). Moreover, the saturation of the transition relation by the rewrite relation gives the effect of weak bisimulation [12].

- I-MSOS supports the use of conventional SOS notation for transitions, whereas its foundations are based on MSOS. Extra conciseness is obtained by allowing strictness annotations in funcon signatures, inspired by the \mathbb{K} Framework [9, page 9], reflecting the distinction between values and computations. Strictness annotations in CBS are interpreted as standard congruence rules for the transition relation, avoiding the need for the reversible ‘heating’ and ‘cooling’ rewrites used in \mathbb{K} .

2.2. The PLANCOMPS Project

The CBS framework has been developed by PLANCOMPS. This project was originally funded by EPSRC in the UK (2011–16), and it is currently continuing as an open international collaboration.

Precursors of CBS include an embedding of component-based semantics in PROLOG: language syntax and its translation to funcons were specified together using Definite Clause Grammars; MSOS transition rules were written as inference rules, supporting the notation introduced in [4] for transitions with labels, and their desugaring into PROLOG rules was also implemented in PROLOG. A similar system, MMT [13], was implemented using the MAUDE framework instead of PROLOG. Both systems have been used in connection with teaching operational semantics [14].

PLANCOMPS developed the PROLOG-based meta-notation further to support value-computation rules and rewrites, as well as typing rules for static semantics. It was used to validate the component-based semantics of CAML LIGHT [5]. PLANCOMPS subsequently switched from PROLOG to HASKELL as implementation language, as described in Sect. 6 and [15].

2.3. Other Semantic Frameworks

Executable semantic definitions have been developed since the early 1970s. Current frameworks supporting executable semantics² include the COREASM tools [16], OTT [17], PLT REDEX [18], the \mathbb{K} tools [19], MAUDE [20] and MMT [13], MELANGE [21], and DYNSEM [22]. Most of these frameworks have a high degree of modularity. An extensive discussion relating several of these other frameworks to component-based semantics can be found in [5].

Apart from CBS, the only current framework that comes together with an extensive collection of reusable components is MMT, which is based on an embedding of MSOS in MAUDE. Case studies using MMT include a sublanguage of CONCURRENT ML, and MINIJAVA. However, MMT was developed before I-MSOS and value-computation systems, and lack of support for these features in MMT significantly reduces the conciseness and readability of its funcon definitions.

It is of considerable interest to develop executable definitions of collections of funcons in some of the other frameworks listed above: that would allow users of those frameworks to specify languages simply by translation to funcons, as in CBS; and it would allow CBS users to exploit features of those frameworks that have not been implemented in CBS.

We have already investigated the possibility of defining funcons in \mathbb{K} , and of using the \mathbb{K} tools to translate languages to funcons [23]. Similarly to CBS, the \mathbb{K} rules for funcons do not require modification when funcons are combined or new funcons added. However, in contrast to CBS, the \mathbb{K} Framework requires configurations of entities to be defined monolithically. Different collections of funcons may require different sets of entities, so each translation of a language to funcons needs to be accompanied by a configuration definition. Any reuse of parts of configurations between different language definitions is left implicit.

The DYNSEM meta-language [22] has been used³ to specify a semantics for a dynamic demonstration language called SIMPLELANGUAGE.⁴ One of the present authors has developed a translation of (most of)

²We restrict attention to formal semantic frameworks with established foundations.

³<https://github.com/MetaBorgCube/metaborg-s1>

⁴<https://github.com/graalvm/simplelanguage>

```

Pgm : pgm      ::= decls
Decls : decls  ::= decl decls?
Decl : decl    ::= func-decl | vars-decl
FuncDecl : func-decl ::= 'function' id '(' ids? ')' block
Ids : ids      ::= id (',' ids)?
VarsDecl : vars-decl ::= 'var' declarators ';'
Declarators : declarators ::= declarator (',' declarators)?
Declarator : declarator ::= id | id '=' exp | id ranks
Ranks : ranks  ::= '[' exps ']' ranks?
Block : block  ::= '{' stmts? '}'
Stmts : stmts  ::= stmt stmts?
Stmt : stmt    ::= imp-stmt | vars-decl
ImpStmt : imp-stmt ::= block
                        | exp ';'
                        | 'if' '(' exp ')' block ('else' block)?
                        | 'while' '(' exp ')' block
                        | 'for' '(' stmt exp ';' exp ')' block
                        | 'print' '(' exps ')' ';'
                        | 'return' exp? ';'
                        | 'try' block 'catch' '(' id ')' block
                        | 'throw' exp ';'
Exp : exp      ::= '(' exp ')' | value | lexp
                        | lexp '=' exp | '++' lexp
                        | '-' exp | exp '(' exps? ')'
                        | 'sizeof' '(' exp ')' | 'read' '(' ')'
                        | exp '+' exp | exp '-' exp
                        | exp '*' exp | exp '/' exp | exp '%' exp
                        | exp '<' exp | exp '<=' exp
                        | exp '>' exp | exp '>=' exp
                        | exp '==' exp | exp '!=' exp
                        | '!' exp | exp '&&' exp | exp '||' exp
LExp : lexp    ::= id | lexp '[' exps ']'
Exps : exps    ::= exp (',' exps)?
V : value     ::= bool | int | string

```

Figure 1: CBS specification of SIMPLE abstract syntax (omitting concurrent constructs).

SIMPLELANGUAGE to funcons in CBS (available online [7]), and experimented with defining those funcons in DYNSEM. DYNSEM has adopted implicit propagation of semantic entities from I-MSOS. The DYNSEM definitions of funcons appear to be just as modular as in CBS. However, DYNSEM is based on big-step rules, which do not support specification of interleaving or concurrency.

3. Specifying Programming Languages in CBS

A component-based semantics for a programming language is a compositional translation from the abstract syntax of the language to funcon terms. CBS provides notation for context-free grammars defining abstract syntax, and notation for introducing translation functions and defining them inductively.

As a running example, we will use the SIMPLE programming language, a small C-like imperative language. SIMPLE was introduced by Roşu and Şerbănuţă [9] as a pedagogical language for demonstrating the \mathbb{K} Semantic Framework [19]. This section presents illustrative extracts from our CBS definition of SIMPLE, which is listed in Appendix A. For more detailed discussion and examples of the translation of programming languages to funcons (though not using CBS), see [5] or [23].

3.1. Abstract Syntax Definitions

Our CBS definition of the abstract syntax of SIMPLE is presented in Figure 1. We typeset non-terminal symbols in **sans-serif** and variables in *Capitalised Italic*. Terminal symbols consist of concrete syntax, and are typeset in **'bold'** and enclosed in single quotation marks. A '?' suffix applied to a symbol (or parenthesised sequence of symbols) denotes that the symbol(s) are optional. Each non-terminal may be declared with an associated variable, which can then be used in translation equations to match phrases of that non-terminal sort.

3.2. Translation Functions

We specify the semantics of a programming language by defining a set of compositional *translation functions*. A translation function comprises a signature and a set of translation equations. The signature declares the sort of source language phrases that the function translates, and the sort of funcon term that it produces. For example, the signature of a translation function named *val* that translates SIMPLE **value** phrases into funcon terms computing **values** is declared as follows:

$$\textit{Semantics } \mathit{val}[_ : \mathit{value}] : \Rightarrow \mathbf{values} \quad (1)$$

The convention in the funcon framework is to use the plural form for funcon type names. We typeset funcons and their types in **bold**, and translation functions in *italic*.

In SIMPLE, the semantics of an expression depends on whether it occurs to the left or right of an assignment operator (so-called *l-expressions* and *r-expressions*). An l-expression should evaluate to an imperative variable, whereas an r-expression, if it results in a variable, should evaluate to the value currently assigned to that variable. In the grammar we have grouped the subset of expressions that are valid l-expressions into a separate non-terminal named **lexp**, and we define two translation functions for specifying the semantics of expressions:

$$\textit{Semantics } \mathit{lval}[_ : \mathit{lexp}] : \Rightarrow \mathbf{variables} \quad (2)$$

$$\textit{Semantics } \mathit{rval}[_ : \mathit{exp}] : \Rightarrow \mathbf{values} \quad (3)$$

A signature is accompanied by a set of translation equations. Each equation should be *compositional*: translation functions in the right-hand side of an equation should be applied to subphrases of the argument phrase in the left-hand side of the equation, for example:

$$\textit{Rule } \mathit{rval}[\mathbf{'!'} \mathit{Exp}] = \mathbf{not}(\mathit{rval}[\mathit{Exp}]) \quad (4)$$

$$\textit{Rule } \mathit{rval}[\mathit{Exp}_1 \mathbf{ '+' } \mathit{Exp}_2] = \mathbf{integer-add}(\mathit{rval}[\mathit{Exp}_1], \mathit{rval}[\mathit{Exp}_2]) \quad (5)$$

$$\textit{Rule } \mathit{rval}[\mathbf{'read' ' (' ')}] = \mathbf{read} \quad (6)$$

The translation of an **lexp** phrase that occurs as an r-expression is:

$$\textit{Rule } \mathit{rval}[\mathit{LExp}] = \mathbf{assigned}(\mathit{lval}[\mathit{LExp}]) \quad (7)$$

where:

$$\textit{Rule } \mathit{lval}[\mathit{Id}] = \mathbf{bound}(\mathit{id}[\mathit{Id}]) \quad (8)$$

$$\textit{Rule } \mathit{lval}[\mathit{LExp} \mathbf{' [' Exp ']}] = \mathbf{checked}(\mathbf{index}(\mathbf{integer-add}(1, \mathit{rval}[\mathit{Exp}])), \mathbf{vector-elements}(\mathit{rval}[\mathit{LExp}]))) \quad (9)$$

Notice that dereferencing imperative variables (Rule 7) and identifier bindings (Rule 8) are implicit in SIMPLE syntax, but are made explicit in the produced funcon term.

Programming languages often contain constructs that are intended to be understood as abbreviations for other constructs in the language: so-called 'syntactic sugar'. CBS allows *desugaring equations* to be defined. For example, a SIMPLE **for**-loop can be desugared into a SIMPLE **while**-loop as follows:

$$\begin{aligned} \textit{Rule } [\mathbf{'for' ' (' Stmt Exp_1 ';' Exp_2 ')} \mathbf{' {' Stmts '}}] : \mathbf{stmt} = \\ [\mathbf{' {' Stmt 'while' ' (' Exp_1 ')} \mathbf{' {' {' Stmts '}} Exp_2 ';' '}} \mathbf{' '}] \end{aligned} \quad (10)$$

This desugaring equation mirrors the \mathbb{K} specification of SIMPLE [9, Page 26], which uses analogous “desugaring macros” to achieve the same effect. However, we deviate slightly from the \mathbb{K} desugaring, which (presumably inadvertently) scopes any declarations in the loop body ($Stmts$) over the for-iterator (Exp_2).

Most of the equations presented so far have translated each SIMPLE construct directly into a corresponding funcon. Translating many other SIMPLE constructs is just as straightforward (see Appendix A). However, often the semantics of a language construct needs to be expressed by a combination of funcons. For example, here are the equations for two forms of SIMPLE declarator:

$$\text{Semantics } \mathit{var-declare} [- : \text{declarator}] : \Rightarrow \text{environments} \quad (11)$$

$$\text{Rule } \mathit{var-declare} [Id] = \mathbf{bind}(id [Id], \mathbf{allocate-variable}(\mathbf{values})) \quad (12)$$

$$\text{Rule } \mathit{var-declare} [Id '=' Exp] = \mathbf{bind}(id [Id], \mathbf{allocate-initialised-variable}(\mathbf{values}, \mathit{rval} [Exp])) \quad (13)$$

One of our aims is that the name of a funcon should be sufficiently indicative of its semantics that a reader should be able to understand the gist of a translation without referring to the funcon definitions.

3.3. Language-Specific Funcons

Although the funcon library contains over a hundred computational funcons (as well as over a hundred value operations), sometimes the semantics of a language construct cannot be directly expressed by a straightforward combination of existing funcons. In SIMPLE, declarators for nested arrays are one such construct. The difficulty is that the rank of an outer array is used to determine the size of each inner array, but the expression computing that array size should only be evaluated once, with the result shared between the inner arrays. Roşu and Şerbănuţă also encounter this issue in their \mathbb{K} specification [9, page 28]. They note that this case cannot be handled by desugaring, as that would duplicate the expression (and any side effects it may contain). Their solution, which they remark is not ideal, is to have their semantic rule for array declarators generate SIMPLE code that uses a **for**-loop to iteratively allocate each sub-array. This code generation happens at run-time, after the size expression has been computed.

We cannot take this approach in the funcon framework: funcon definitions cannot refer to source-language syntax or invoke translation functions. However, CBS provides a convenient alternative: local funcons can be specified within a language definition. This is appropriate for funcons that are not considered reusable enough to be included in the funcon library, but nonetheless capture a useful concept of the language being defined. (Local funcon definitions are also useful for validating new funcons before adding them to the library.) Here, we introduce a SIMPLE-specific funcon **allocate-nested-vectors** that precisely captures the semantics of nested array declarators, allowing us to define the translation as follows:

$$\text{Rule } \mathit{var-declare} [Id Ranks] = \mathbf{bind}(id [Id], \mathbf{allocate-nested-vectors}[ranks [Ranks]]) \quad (14)$$

The definition of **allocate-nested-vectors** is included in Appendix A.4.2.

3.4. Remarks

Our definition of SIMPLE omits the concurrent aspects of the language. This is not due to any limitation of CBS or (I-)MSOS-based specifications – indeed, one author of this article has previously used MSOS to specify a sublanguage of Concurrent ML [24] – rather it is that development of the funcon library is ongoing work, and funcons specifying concurrency have not yet been added.

4. Specifying Fundamental Constructs in CBS

Structural Operational Semantics (SOS) [10] is a well-established framework for specifying computational behaviour, where the behaviour of programs is modelled by labelled transition systems, defined inductively by axioms and inference rules. To specify the dynamic semantics of funcons, we use small-step rules in a *modular* variant of SOS called MSOS [4]. The CBS notation for MSOS rules is based on *Implicitly-Modular SOS* (I-MSOS) [3], which provides a notational style similar to conventional SOS. CBS (or I-MSOS) rules can be understood as syntactic sugar for MSOS rules. We assume the reader is familiar with conventional SOS rules (e.g. [10]). The rules define two transition relations, for computational steps and context-free

rewrites respectively. A formal account of the underlying MSOS framework can be found in [4], of the I-MSOS rules in [3], and of the two transition relations in [11].

This section presents an informal and example-driven overview of using CBS to specify funcons. The funcons in this section are selected to demonstrate different aspects of CBS specifications. Appendix B.1 lists the signatures of all funcons used in the SIMPLE language specification; the definitions of those funcons are available online [7].

4.1. Transition Relations

Consider the following CBS specification of **if-true-else**:

$$\text{Rule } \frac{B \longrightarrow B'}{\mathbf{if-true-else}(B, X, Y) \longrightarrow \mathbf{if-true-else}(B', X, Y)} \quad (15)$$

$$\text{Rule } \mathbf{if-true-else}(\mathbf{true}, X, _) \rightsquigarrow X \quad (16)$$

$$\text{Rule } \mathbf{if-true-else}(\mathbf{false}, _, Y) \rightsquigarrow Y \quad (17)$$

This is a fairly conventional SOS specification of a conditional expression, except that there are two transition relations: ‘ \longrightarrow ’ and ‘ \rightsquigarrow ’. The former denotes exactly one step of computation, which in general may have arbitrary side effects and be sensitive to its context. The latter denotes context-free term rewriting, which cannot involve side effects, and is reflexive and transitive. The distinction facilitates reasoning about funcons terms. It can also be exploited by a funcon interpreter to provide a more efficient implementation of the pure transitions. The formal details of these two relations, and the interaction between them, are spelled out in [11].⁵

4.2. Funcon Signatures

CBS requires each funcon to have a declared signature. For example, the signature for **if-true-else** is declared as follows:

$$\text{Funcon } \mathbf{if-true-else}(_ : \mathbf{booleans}, _ : \Rightarrow T, _ : \Rightarrow T) : \Rightarrow T \quad (18)$$

The result of the funcon, and each funcon parameter, has a *sort*. A funcon sort is either a *value sort* (a type), or a *computation sort* (a type prefixed by the ‘ \Rightarrow ’ symbol, which is pronounced “computes”). When the result sort is a value sort, then the funcon is a value constructor, and the funcon will not have any operational rules.

Additionally, polymorphic funcons may have signatures that contain type variables. In the case of **if-true-else**, the type variable T is used to express that the second and third arguments compute values of the same type, and that the funcon as a whole also computes a value of that type.

Notice that the first parameter has a value sort. This does not mean that **if-true-else** is limited to being applied to literal Boolean values; rather, value-sort annotations play a special role in funcon signatures, as we will now describe.

When a parameter has a value sort, we say that the funcon is *strict* in that argument. If a funcon is strict in all its arguments, then we say that the funcon is strict. For each *strict* parameter in a funcon’s signature, a *congruence rule* is implicitly generated for that argument.

A congruence rule has a single premise consisting of a computation step for the strict argument. The target of the rule’s conclusion is the same as the source of the conclusion, except that the argument subterm is replaced with the updated subterm from the target of the premise. For example, Rule 15 is the congruence rule implicitly generated by the signature of **if-true-else** (18), and thus that rule can be omitted. A funcon with congruence rules generated for several arguments is non-deterministic with respect to the order in which these arguments are evaluated: the congruence rules admit all possible interleavings. This technique for automatically generating congruence rules was inspired by the strictness annotations in the \mathbb{K} Framework [9, page 9].

⁵ In [11] the relations ‘ \rightsquigarrow ’ and ‘ \longrightarrow ’ are written ‘ \Rightarrow ’ and ‘ \rightarrow ’, respectively.

4.3. Values and Types

CBS provides a rich suite of primitive and composite values, and operations on those values. The available operations include arithmetic, conversions between types, and insertion/deletion from collections, among others. The signatures of all value operations used in this article are listed in Appendix B.2. Some familiar types of values are left unspecified, for example **integers**, **floats** and **maps**. In these cases the semantics of value operations is implemented manually in a library provided as part of the Haskell funcon framework (see Sect. 6.5).

The names of types are themselves funcon values, and belong to the type **types**. All values, including type names, are members of the type **values** (i.e. **values** is a supertype of **integers**, **floats**, etc.). The type **values** is frequently used when specifying funcons that can take arbitrary values as arguments.

New data types can be introduced. For example, the data types **booleans**, and **null-type** are defined as follows:

$$\text{Datatype } \mathbf{booleans} ::= \mathbf{true} \mid \mathbf{false} \tag{19}$$

$$\text{Datatype } \mathbf{null-type} ::= \mathbf{null} \tag{20}$$

The type **null-type** is the unit type of the funcon library, and is used as the result type of commands.

When specifying a funcon, it may be necessary to restrict a rule to be only applicable when some of its arguments are values (rather than computations). This can be achieved by annotating the appropriate arguments with ‘: T ’ (or a more specific type).

Alternatively, arguments in rules can be restricted to values by using *patterns*. A pattern consists of a value constructor, with sub-patterns to match the arguments of the constructor (if any). Those sub-patterns may include meta-variables. For example, rules 16 and 17 have the **true** and **false** constructors as the first argument, respectively, restricting each rule to that specific argument.

Following [11], the only funcons permitted in CBS patterns are value constructors. This ensures that the internal structure of a computation cannot be observed: the only way to observe anything about a computation is to execute it. Conversely, CBS values can be observed, but must be stable: values do not perform computation steps.

4.4. Semantic Entities

Formal specifications of programming-language constructs typically make use of auxiliary *semantic entities* to model computational side effects. Common examples of such entities include environments, stores or output sequences.⁶ In an SOS specification, such semantic entities are explicit arguments of transition relations, and thus they must appear in every transition formula of that relation. For example, in a conventional SOS specification, Rule 15 could well have been written as follows:

$$\frac{\rho \vdash \langle B, \sigma \rangle \xrightarrow{\alpha} \langle B', \sigma' \rangle}{\rho \vdash \langle \mathbf{if-true-else}(B, X, Y), \sigma \rangle \xrightarrow{\alpha} \langle \mathbf{if-true-else}(B', X, Y), \sigma' \rangle} \tag{21}$$

ρ , σ , σ' and α are all variables. ρ is the environment in which the term is executing, σ and σ' are the state of the store before and after the transition, and α is the output sequence being emitted. Observe that both the environment and output are the same in both the premise and conclusion. This would typically be understood as propagating the environment unmodified from the conclusion to the premise, and propagating the output unmodified from the premise to the conclusion. The use of σ and σ' would typically be understood as propagating the initial store unmodified from the source of the conclusion to the source of the premise, and propagating the resulting store from the target of the premise to the target of the conclusion. While this rule does not access or modify the store directly, the computation step in the premise may do so.

This systematic propagation of semantic entities is common to many rules in SOS specifications. The key feature of I-MSOS is that any semantic entities that are not mentioned in a rule are *implicitly propagated*

⁶The use of environments and stores in SOS dates back to 1981, when they were adopted by Plotkin [10]. But these were simply first-order versions of the corresponding entities used in denotational semantics since the early 1970s (foreshadowed in Strachey’s work on CPL [25], and his lecture notes from 1967 [26]).

between the premise(s) and conclusion, systematically following the propagation schemes described above. This allows semantic entities that are not affected by the programming construct being specified to be omitted from the rule. Thus, rather than Rule 21, in CBS we are able to write the much more concise Rule 15.

Moreover, beyond just notational convenience, implicit propagation brings a key modularity benefit: new semantic entities can later be added without requiring any change to the previously given rules. This high degree of modularity is what allows the collection of funcons to be open-ended, while still ensuring that once a funcon is added to the collection, its CBS definition need never be updated.

To allow semantic entities to be implicitly propagated, they must be classified according to their desired propagation scheme. CBS supports five classes of semantic entity, each with its own notation and system of implicit propagation: *contextual*, *mutable*, *input*, *output* and *control-flow*. The input and control-flow entities are novel: in prior publications on MSOS or funcons (e.g. [4, 5, 24, 27]), only contextual, mutable and output entities were provided, a more limited form of control-flow was achieved by using output entities, and program input was represented by mutable entities.⁷

CBS allows multiple entities of each class, so each entity is tagged with a distinguishing name. The notation for a CBS computation step using one entity of each of the five classes is as follows (writing the classification names as placeholders for the actual entity names):

$$\text{Rule } \text{contextual}(_) \vdash \langle _ , \text{mutable}(_) \rangle \xrightarrow{\text{input?}(_) \text{ control-flow}(_) \text{ output!}(_)} \langle _ , \text{mutable}(_) \rangle \quad (22)$$

Note that the ‘ \rightsquigarrow ’ relation never involves semantic entities.

We will now present examples of each class of semantic entity, and of funcons whose semantics require them.

4.4.1. Contextual Entities

Contextual entities model contextual information that is available to a computation step, but which cannot be modified by that computation step. In a CBS rule, unmentioned contextual entities are implicitly propagated from the transition in the conclusion to the transition in any premises.

The classic example of a contextual entity is an *environment* containing bindings for program identifiers. We declare such an entity in CBS as follows:

$$\text{Entity } \text{environment}(_ : \text{environments}) \vdash _ \longrightarrow _ \quad (23)$$

Here, **environment** is the name of the entity being declared, **environments** is the type of value contained in the entity. The type **environments** is a mapping from **identifiers** to optional values:

$$\text{Funcon } \text{environments} : \Rightarrow \text{types} \quad (24)$$

$$\text{Rule } \text{environments} \rightsquigarrow \text{maps}(\text{identifiers}, \text{values}^?) \quad (25)$$

A ‘?’ suffix on a funcon type (or type variable) denotes an option type; that is, zero or one values of that type. An absent value in an environment represents that a binding is present, but inaccessible.⁸

Let us consider some funcons that use this entity. The funcon **bound**(*I*) evaluates to the value bound to the identifier *I* in the current environment. A common use of this funcon is when specifying the translation of occurrences of identifiers in expressions (e.g. in Eq. 8). The definition of **bound** for non-recursive bindings is as follows:⁹

⁷In prior publications, *contextual*, *mutable* and *output* entities have appeared under a variety of names, most commonly *read-only*, *read-write* and *write-only*, respectively.

⁸For example, in C^\sharp a local variable is in scope in its entire enclosing block (shadowing any outer variable of the same name), but it is not accessible until after its point of declaration.

⁹See [7] for the generalisation to recursive bindings.

$$\text{Funcon } \mathbf{bound}(_ : \mathbf{identifiers}) : \Rightarrow \mathbf{values} \quad (26)$$

$$\text{Rule } \frac{\mathbf{lookup}(\rho, I) \rightsquigarrow V : \mathbf{values}}{\mathbf{environment}(\rho) \vdash \mathbf{bound}(I : \mathbf{identifiers}) \longrightarrow V} \quad (27)$$

$$\text{Rule } \frac{\mathbf{lookup}(\rho, I) \rightsquigarrow ()}{\mathbf{environment}(\rho) \vdash \mathbf{bound}(I : \mathbf{identifiers}) \longrightarrow \mathbf{fail}} \quad (28)$$

The conclusions of Rules 27 and 28 are computation steps, not rewrites, because they refer to a semantic entity. The notation ‘ $()$ ’ denotes the absence of a CBS value.

The funcon $\mathbf{scope}(\rho, X)$ adds the bindings in ρ to the current environment for the purposes of executing the subterm X :

$$\text{Funcon } \mathbf{scope}(_ : \mathbf{environments}, _ : \Rightarrow T) : \Rightarrow T \quad (29)$$

$$\text{Rule } \frac{\mathbf{map-override}(\rho_1, \rho_0) \rightsquigarrow \rho_2 \quad \mathbf{environment}(\rho_2) \vdash X \longrightarrow X'}{\mathbf{environment}(\rho_0) \vdash \mathbf{scope}(\rho_1 : \mathbf{environments}, X) \longrightarrow \mathbf{scope}(\rho_1, X')} \quad (30)$$

$$\text{Rule } \mathbf{scope}(_ : \mathbf{environments}, V : T) \rightsquigarrow V \quad (31)$$

Recall that a funcon signature containing any strict arguments implicitly generates congruence rules for evaluating those strict arguments. Thus the first argument to \mathbf{scope} could be a funcon term that computes an environment. This is indeed the case in SIMPLE, where declarations involve allocating and initialising imperative variables (see Sect. 3.2).

Another example of a contextual entity is $\mathbf{given-value}$:

$$\text{Entity } \mathbf{given-value}(_ : \mathbf{values}^?) \vdash _ \longrightarrow _ \quad (32)$$

This entity is used for a very simple form of binding and scoping. It holds only a single value. The use of $\mathbf{given-value}$ is sufficiently common that we give it special treatment in funcon signatures. A sort $S \Rightarrow T$ represents a term that, in a context where the given value has type S , computes a value of type T . To access the value, we use the funcon \mathbf{given} , which is analogous to \mathbf{bound} :

$$\text{Funcon } \mathbf{given} : T \Rightarrow T \quad (33)$$

$$\text{Rule } \mathbf{given-value}(V : \mathbf{values}) \vdash \mathbf{given} \longrightarrow V \quad (34)$$

$$\text{Rule } \mathbf{given-value}() \vdash \mathbf{given} \longrightarrow \mathbf{fail} \quad (35)$$

To provide a value to a subterm via the $\mathbf{given-value}$ entity, we use the funcon \mathbf{give} , which is analogous to \mathbf{scope} :

$$\text{Funcon } \mathbf{give}(_ : S, _ : S \Rightarrow T) : \Rightarrow T \quad (36)$$

$$\text{Rule } \frac{\mathbf{given-value}(V) \vdash X \longrightarrow X'}{\mathbf{given-value}(_) \vdash \mathbf{give}(V : S, X) \longrightarrow \mathbf{give}(V, X')} \quad (37)$$

$$\text{Rule } \mathbf{give}(_ : S, V : T) \rightsquigarrow V \quad (38)$$

A common use of \mathbf{give} and \mathbf{given} is to model passing arguments to a function, as we shall discuss in Sect. 4.5.

4.4.2. Mutable Entities

Mutable entities model persistent information that can be modified by a computation step: the resulting value of a mutable entity after a step is the initial value of the mutable entity for the next step. In a CBS rule, unmentioned mutable entities are implicitly propagated unmodified.

The classic example of a mutable entity is a *store* containing the values of imperative variables. We declare such an entity in CBS as follows:

$$\text{Entity } \langle _, \mathbf{store}(_ : \mathbf{stores}) \rangle \longrightarrow \langle _, \mathbf{store}(_ : \mathbf{stores}) \rangle \quad (39)$$

The type **stores** is a mapping from locations (**locations**) to optional values, where a mapping to an absent value represents a location without an assignment.

$$\text{Funcon } \mathbf{stores} : \Rightarrow \mathbf{types} \quad (40)$$

$$\text{Rule } \mathbf{stores} \rightsquigarrow \mathbf{maps}(\mathbf{locations}, \mathbf{values}^?) \quad (41)$$

A variable wraps a location and the type of values that can be assigned to that location.

$$\text{Datatype } \mathbf{variables} ::= \mathbf{variable}(_ : \mathbf{locations}, _ : \mathbf{types}) \quad (42)$$

To assign a value to a variable in the **store**, we use the funcon **assign**:

$$\text{Funcon } \mathbf{assign}(_ : \mathbf{variables}, _ : \mathbf{values}) : \Rightarrow \mathbf{null-type} \quad (43)$$

$$\text{Rule } \frac{\mathbf{and}(\mathbf{is-in-set}(L, \mathbf{dom}(\sigma)), \mathbf{is-in-type}(V, T)) \rightsquigarrow \mathbf{true} \quad \mathbf{map-override}(\{L \mapsto V\}, \sigma) \rightsquigarrow \sigma'}{\langle \mathbf{assign}(\mathbf{variable}(L : \mathbf{locations}, T : \mathbf{types}), V : \mathbf{values}), \mathbf{store}(\sigma) \rangle \longrightarrow \langle \mathbf{null}, \mathbf{store}(\sigma') \rangle} \quad (44)$$

$$\text{Rule } \frac{\mathbf{and}(\mathbf{is-in-set}(L, \mathbf{dom}(\sigma)), \mathbf{is-in-type}(V, T)) \rightsquigarrow \mathbf{false}}{\langle \mathbf{assign}(\mathbf{variable}(L : \mathbf{locations}, T : \mathbf{types}), V : \mathbf{values}), \mathbf{store}(\sigma) \rangle \longrightarrow \langle \mathbf{fail}, \mathbf{store}(\sigma) \rangle} \quad (45)$$

The notation $\{ _ \mapsto _ \}$ constructs map values. The value assigned to a variable is retrieved using the funcon **assigned** (e.g. in Eq. 7):

$$\text{Funcon } \mathbf{assigned}(_ : \mathbf{variables}) : \Rightarrow \mathbf{values} \quad (46)$$

$$\text{Rule } \frac{\mathbf{lookup}(\sigma, L) \rightsquigarrow V : \mathbf{values}}{\langle \mathbf{assigned}(\mathbf{variable}(L : \mathbf{locations}, T : \mathbf{types})), \mathbf{store}(\sigma) \rangle \longrightarrow \langle V, \mathbf{store}(\sigma) \rangle} \quad (47)$$

$$\text{Rule } \frac{\mathbf{lookup}(\sigma, L) \rightsquigarrow ()}{\langle \mathbf{assigned}(\mathbf{variable}(L : \mathbf{locations}, T : \mathbf{types})), \mathbf{store}(\sigma) \rangle \longrightarrow \langle \mathbf{fail}, \mathbf{store}(\sigma) \rangle} \quad (48)$$

4.4.3. Output Entities

Output entities model emitted output. An output entity is a sequence of values, and the sequences from adjacent computation steps are implicitly concatenated, with non-terminating computations potentially resulting in infinite output. If an output entity is not mentioned in a rule, it is implicitly propagated from the premise (if any) to the conclusion.

An example of an output entity is the standard-output channel. We declare output entities with an ‘!’ suffix:

$$\text{Entity } _ \xrightarrow{\mathbf{standard-out}!(_ : \mathbf{values}^*)} _ \quad (49)$$

The **standard-out** entity holds a sequence of values. A funcon for printing a value to the **standard-out** entity can be defined as follows:

$$\text{Funcon } \mathbf{print}(_ : \mathbf{values}^*) : \Rightarrow \mathbf{null-type} \quad (50)$$

$$\text{Rule } \mathbf{print}(V^* : \mathbf{values}^*) \xrightarrow{\mathbf{standard-out}!(V^*)} \mathbf{null} \quad (51)$$

The star in ‘ V^* ’ and the annotation ‘ \mathbf{values}^* ’ denote that ‘ V^* ’ is a placeholder for zero or more **values**.

4.4.4. Input Entities

Input entities model consumed input. Like output entities they form a sequence of values, but this represents the input *consumed* by a computation step, not the output emitted.

An example of an input entity is the standard-input channel. We declare input entities with a ‘?’ suffix:

$$\text{Entity } _ \xrightarrow{\mathbf{standard-in}?(_ : \mathbf{values}^*)} _ \quad (52)$$

The **standard-in** entity holds a sequence of values, with **null** representing the end of input. A funcon for reading a single value from the **standard-in** entity can be defined as follows:

$$\text{Funcon } \mathbf{read} : \Rightarrow \mathbf{values} \quad (53)$$

$$\text{Rule } \mathbf{read} \xrightarrow{\text{standard-in?}(V : \sim \mathbf{null-type})} V \quad (54)$$

$$\text{Rule } \mathbf{read} \xrightarrow{\text{standard-in?}(\mathbf{null})} \mathbf{fail} \quad (55)$$

The notation ‘ \sim ’ is type complement; that is, Rule 54 is applicable when the input value is not **null**.

The Haskell funcon framework discussed in Sect. 6 supports simulated and interactive input, where the input values are provided before execution begins, or given by the user during execution, respectively.

4.4.5. Control-flow Entities

Control-flow entities model changes in the control-flow of the program. In CBS rules, control-flow entities contain signals that communicate between two rules in a derivation tree. Propagation of control-flow entities is not *directed*; rather, any signals in unmentioned control-flow entities are implicitly synchronised between the premise and conclusion of a rule. Unlike input and output entities, a control-flow entity does not accumulate a sequence of values. Instead, it contains either a present signal, which carries a value, or an absent signal, which is empty.

The classic example of a control-flow change is abrupt termination of a computation. The **abrupted** control-flow entity indicates abrupt termination when it appears with value V in a label on a transition. The type of value is used to indicate different forms of abrupt termination, which include thrown exceptions, returning from a function, or breaking from a loop. The absence of V indicates the absence of abrupt termination.

$$\text{Entity } _ \xrightarrow{\text{abrupted}(_ : \mathbf{values}?) } _ \quad (56)$$

Abrupt termination is initiated using **abrupt**, which emits a signal in the **abrupted** entity:

$$\text{Funcon } \mathbf{abrupt}(_ : \mathbf{values}) : \Rightarrow \mathbf{empty-type} \quad (57)$$

$$\text{Rule } \mathbf{abrupt}(V : \mathbf{values}) \xrightarrow{\text{abrupted}(V)} \mathbf{stuck} \quad (58)$$

Funcons for throwing exceptions, for example, are then defined as follows:

$$\text{Datatype } \mathbf{throwing} ::= \mathbf{thrown}(_ : \mathbf{values}) \quad (59)$$

$$\text{Funcon } \mathbf{throw}(_ : \mathbf{values}) : \Rightarrow \mathbf{empty-type} \quad (60)$$

$$\text{Rule } \mathbf{throw}(V : \mathbf{values}) \rightsquigarrow \mathbf{abrupt}(\mathbf{thrown}(V)) \quad (61)$$

To handle thrown exceptions, we use a funcon called **handle-throw**:

$$\text{Funcon } \mathbf{handle-throw}(_ : S \Rightarrow T, _ : X \Rightarrow T) : S \Rightarrow T \quad (62)$$

$$\text{Rule } \frac{X \xrightarrow{\text{abrupted}(\mathbf{thrown}(V))} _}{\mathbf{handle-throw}(X, H) \xrightarrow{\text{abrupted}()} \mathbf{give}(V, H)} \quad (63)$$

$$\text{Rule } \frac{X \xrightarrow{\text{abrupted}(V : \sim \mathbf{throwing})} X'}{\mathbf{handle-throw}(X, H) \xrightarrow{\text{abrupted}(V)} \mathbf{handle-throw}(X', H)} \quad (64)$$

$$\text{Rule } \frac{X \xrightarrow{\text{abrupted}()} X'}{\mathbf{handle-throw}(X, H) \xrightarrow{\text{abrupted}()} \mathbf{handle-throw}(X', H)} \quad (65)$$

$$\text{Rule } \mathbf{handle-throw}(V : T, _) \rightsquigarrow V \quad (66)$$

The first argument of **handle-throw** is a term to be evaluated, and the second argument is a term that serves as an exception handler. If the step on the first argument emits a **thrown** signal with value V , Rule 63 invokes the handler H by giving it V . Other forms of abrupt termination are not handled, and are

propagated by Rule 64. Rule 65 is similar to a congruence rule for the first argument, except that the step on the argument must not emit a **thrown** signal. If the first argument of **handle-thrown** is a value, Rule 66 ignores the handler and returns that value.

As another example, a funcon representing returning from within a function is defined as follows:

$$\text{Datatype } \mathbf{returning} ::= \mathbf{returned}(_ : \mathbf{values}) \quad (67)$$

$$\text{Funcon } \mathbf{return}(_ : T) : \Rightarrow \mathbf{empty-type} \quad (68)$$

$$\text{Rule } \mathbf{return}(V : T) \rightsquigarrow \mathbf{abrupt}(\mathbf{returned}(V)) \quad (69)$$

We then define a funcon **handle-return** that ‘handles’ returned values by evaluating its argument and gives V if its argument terminates normally with value V , or abruptly for reason **returned**(V). This is analogous to **handle-thrown**, so we omit the rules. Further examples of using control-flow entities are in Sect. 5.

4.5. Functions and Closures

Semantic entities hold *values*, not arbitrary computation terms. However, sometimes we want to record an unevaluated computation; for example, we need to represent procedural abstractions as values in the environment.

The value constructor **abstraction** takes a computation as argument and constructs a value:

$$\text{Funcon } \mathbf{abstraction}(_ : S \Rightarrow T) : \mathbf{abstractions}(S \Rightarrow T) \quad (70)$$

Procedural abstractions are modelled using the **functions** type defined below.

$$\text{Datatype } \mathbf{functions}(S, T) ::= \mathbf{function}(_ : \mathbf{abstractions}(S \Rightarrow T)) \quad (71)$$

We do not merge **functions** and **abstractions** into a single funcon, because we also model other concepts using abstractions, including patterns (which compute bindings), thunks (delayed computations), and continuations (which we discuss in Sect. 5).

We model argument passing using the **given-value** entity: **given** refers to the argument (which could be a tuple of arguments) inside the function body at the definition site of the function, and **give** supplies the argument at the call site of the function. Function application is modelled by extracting the computation from the **function** value, and *giving* an argument to that computation:

$$\text{Funcon } \mathbf{apply}(_ : \mathbf{functions}(S, T), _ : S) : \Rightarrow T \quad (72)$$

$$\text{Rule } \mathbf{apply}(\mathbf{function}(\mathbf{abstraction}(X)), V : S) \rightsquigarrow \mathbf{give}(V, X) \quad (73)$$

Note that we are *not* saying that we translate occurrences of SIMPLE identifiers in a function body to the funcon **given**. Rather, we use **given** to refer to the argument when expressing the semantics of parameter binding. The bindings are then scoped over the function body, using **scope**. (See Appendix A.4.3.)

This treatment of functions leads to *dynamic scoping* for identifiers, as the **bound** funcon is defined such that it looks up the identifier in the environment where it is evaluated. To support *static scoping* of identifiers we use the funcon **closed**, which ensures that its computation argument is evaluated without any non-local bindings:

$$\text{Funcon } \mathbf{closed}(_ : \Rightarrow T) : \Rightarrow T \quad (74)$$

$$\text{Rule } \frac{\mathbf{environment}(\mathbf{map-empty}) \vdash X \longrightarrow X'}{\mathbf{environment}(_) \vdash \mathbf{closed}(X) \longrightarrow \mathbf{closed}(X')} \quad (75)$$

$$\text{Rule } \mathbf{closed}(V : T) \rightsquigarrow V \quad (76)$$

The closure of a computation X is an abstraction with a closed body derived from X and with the current environment available as local bindings:

$$\text{Funcon } \mathbf{closure}(_ : S \Rightarrow T) : \Rightarrow \mathbf{abstractions}(S \Rightarrow T) \quad (77)$$

$$\text{Rule } \mathbf{environment}(\rho) \vdash \mathbf{closure}(X) \longrightarrow \mathbf{abstraction}(\mathbf{closed}(\mathbf{scope}(\rho, X))) \quad (78)$$

The computation X will thus be evaluated with respect to the environment where the closure is formed, not the environment where the closure is evaluated.

4.6. Remarks

We believe language definitions written with CBS are relatively easy to comprehend and accessible to a wide audience. The process of defining the syntax of a language with (a variant of Extended) BNF is ubiquitous and well-understood. Funcon translations are defined through simple translation equations familiar especially to users experienced in denotational semantics. The I-MSOS rules that define funcons are similar to standard SOS rules, with the difference that auxiliary semantic entities may be omitted, making specifications more concise. Moreover, for understanding the behaviour of an individual funcon it is unnecessary to know the details of the implicit propagation rules of the entity classes. Most importantly, to define languages in CBS, awareness of the precise formal semantics of the used funcons is not required: an informal explanation suffices, which is provided in the form of inline documentation (available online [7]).

The library of reusable funcons contains funcons that capture many types of constructs found in programming languages across language paradigms. For example, the funcons give semantics to different evaluation strategies in procedural and functional programming, algebraic data types and patterns in functional programming, classes and inheritance in object-oriented programming, as well as GOTO-statements and other forms of abnormal control flow.

Funcons have not yet been developed for all aspects of existing programming languages. For example, the library of funcons currently contains funcons for interleaving (and its prevention), and for interactive input and output, but not for threads, synchronisation, or distributed processes. Adding appropriate entities for use in CBS rules specifying the semantics of such constructs should be straightforward, based on previous experience with using MSOS for Concurrent ML [24] and on the many examples of process algebra specifications provided in the literature (the latter relies on the close relationship between small-step MSOS and the conventional framework of transition system specifications). Deciding whether to include funcons for synchronous or asynchronous communication (or both) may be more difficult, but case studies can be carried out to inform such choices. The main difficulties seem likely to arise in connection with executability and validation: interpreter generation will need to ensure that CBS rules are applied fairly, and model checking may be needed to determine whether all intended computations are possible. The \mathbb{K} framework addressed those issues [9] by annotating transition rules (essentially identifying different states); it is unclear whether we will be able to adopt a similar approach for interpreter generation based on CBS rules.

In general, if funcons for particular language constructs cannot be defined in CBS as it is presented in this paper, we may extend the current collection of entity classes. Adding new entities would not affect the definitions of existing funcons, assuming that the new entities are implicitly propagated when not mentioned in rules.

We have only considered the *dynamic* semantics of funcons here. Together with Churchill and Torrini, two of the present authors have previously explored specifying the *static* semantics of funcons [5]. Following that approach, we intend to add facilities for declaring typing rules in CBS, and to generate static analyses from those rules; but this remains as future work.

5. Case Study: Delimited Continuations

In Sect. 3, we discussed using CBS to specify the semantics of an existing language (SIMPLE). In this section, we demonstrate the ease with which an existing CBS language specification can be extended, using the addition of a *delimited control operator* as a small case study. We also use this case study to demonstrate how advanced forms of control flow can be specified in CBS.

5.1. Extending SIMPLE with Delimited Continuations

First we extend SIMPLE with syntax for two new forms of statement:

```
ImpStmt : imp-stmt ::= ... | 'shift' id block | 'resume' id ';
```

Informally, the desired semantics of **'shift *Id* *Block*'** is that the remainder of the enclosing block is captured as a *delimited continuation*, bound to *Id*, and scoped over *Block*. Then, *Block* is executed *instead of the*


```

function main() {
  print (1);
  {
    print (2);
    shift k { resume k; print (3); resume k; }
    print (4);
  }
  print (5);
}

```

Figure 2: A program that prints the sequence ‘1 2 4 3 4 5’.

remainder of the enclosing block. During execution of *Block*, the captured continuation can be executed zero or more times, which is achieved by invoking ‘**resume** *Id*’. This **shift** statement is essentially Danvy and Filinski’s [28] *shift* operator, except formulated as a statement rather than an expression. See Fig. 2 for an example.

The funcon repository contains several funcons for expressing delimited control (we defer the accompanying rules to Sect 5.2):

$$\text{Funcon } \mathbf{control}(- : \mathbf{functions}(\mathbf{continuations}(T, S), S)) : \Rightarrow T \quad (79)$$

$$\text{Funcon } \mathbf{delimit-cc}(- : \Rightarrow T) : \Rightarrow T \quad (80)$$

$$\text{Funcon } \mathbf{resume-continuation}(- : \mathbf{continuations}(S, T), - : S) : \Rightarrow T \quad (81)$$

The **control** and **delimit-cc** funcons correspond closely to Felleisen’s [29] *control* and *prompt* operators. The only difference is that captured continuations have type **continuations**, which is distinct from the type **functions**. The **resume-continuation** funcon is analogous to **apply** (Rules 72 & 73).

We formally define the semantics of **shift** and **resume** by translation to funcons. In (extended) SIMPLE, blocks act as delimiters when determining the scope of a captured continuation. The existing translation of blocks (from Appendix A.3) is simply to execute the statements within the block:

$$\text{Rule } \mathit{exec} \llbracket \{ 'Stmts' \} \rrbracket = \mathit{exec} \llbracket Stmts \rrbracket \quad (82)$$

We therefore modify this equation to also insert a continuation delimiter:

$$\text{Rule } \mathit{exec} \llbracket \{ 'Stmts' \} \rrbracket = \mathbf{delimit-cc}(\mathit{exec} \llbracket Stmts \rrbracket) \quad (83)$$

The **control** funcon captures the current continuation in a similar manner to SIMPLE’s **shift** statement. However, SIMPLE’s **shift** statement has an explicit binder as part of the syntax, whereas **control** takes a function as its argument and applies that function to the captured continuation. To bridge this mismatch, the translation of a **shift** statement to funcons explicitly constructs a function that binds the argument (the continuation) to the identifier:

$$\begin{aligned} \text{Rule } \mathit{exec} \llbracket \mathbf{shift} \ 'Id' \ Block \rrbracket = \\ \mathbf{control}(\mathbf{function}(\mathbf{closure}(\mathbf{scope}(\mathbf{bind}(id \llbracket Id \rrbracket), \mathbf{given}), \mathit{exec} \llbracket Block \rrbracket)))) \end{aligned} \quad (84)$$

The signature of the **control** funcon indicates that it computes a value, and thus, in general, the captured continuation may include an enclosing expression. Consequently, the **resume-continuation** funcon requires a value to plug the hole in the expression where **control** occurred. Conversely, SIMPLE’s **shift** is a *statement*, and thus does not require a value to resume. To bridge this mismatch, the translation of a **resume** statement inserts a **null** value:

$$\text{Rule } \mathit{exec} \llbracket \mathbf{resume} \ 'Id' \ ; \rrbracket = \mathbf{resume-continuation}(\mathbf{bound}(id \llbracket Id \rrbracket), \mathbf{null}) \quad (85)$$

This serves as a representative example of how the definition of a language feature often proceeds. The funcon repository does not provide a funcon for every possible variation of every possible language feature.

Rather, it provides funcons that embody core concepts, and then the translation expresses the semantics of the language feature as a combination of those funcons.

This completes the addition of delimited continuations to SIMPLE. We only needed to declare the new syntax, add translation equations for that syntax, and modify one existing translation equation. A cynical reader may consider that we have ‘cheated’: the complexity of the semantics of **shift** is in the definition of **control**, **delimit-cc**, and **resume-continuation**, and we have not included these definitions. But this is precisely our aim: by providing reusable components with pre-specified semantics, the effort required to formally specify a programming language is significantly reduced by moving the majority of the specification work into those reusable components.

5.2. Defining Delimited Control Funcons

We now present the definitions of the funcons that we used to specify the semantics of SIMPLE’s delimited continuations. These are based on the I-MSOS definitions of similar funcons previously presented by two authors of this article, together with Torrini [30]. However, the control-flow entities provided by CBS facilitate a simpler specification than was possible in that prior work.

First we define the **continuations** type in terms of **abstractions**, in a similar manner to **functions** (Rule 71):

$$\text{Datatype } \mathbf{continuations}(S, T) ::= \mathbf{continuation}(_ : \mathbf{abstractions}(() \Rightarrow T)) \quad (86)$$

Next we need to express how **control** captures the current continuation. The key idea is to use control-flow entities to communicate between **control** and its enclosing **delimit-cc**, in a similar manner to how abrupt termination is modelled (Sect. 4.4.5). We define **control** to emit a signal containing the function that is to be applied to the continuation, and leave a **hole** in the term:

$$\text{Rule } \mathbf{control}(F : \mathbf{functions}(_, _)) \xrightarrow{\mathbf{control-signal}(F)} \mathbf{hole} \quad (87)$$

The **delimit-cc** funcon receives the **control-signal**, constructs the continuation, and applies the received function to that continuation:

$$\text{Rule } \frac{X \xrightarrow{\mathbf{control-signal}(F)} X'}{\mathbf{delimit-cc}(X) \xrightarrow{\mathbf{control-signal}()} \mathbf{delimit-cc}(\mathbf{apply}(F, \mathbf{continuation}(\mathbf{closure}(X'))))} \quad (88)$$

We also need rules for when the body of **delimit-cc** does not emit a **control-signal**, and for when it has been evaluated to a value:

$$\text{Rule } \frac{X \xrightarrow{\mathbf{control-signal}()} X'}{\mathbf{delimit-cc}(X) \xrightarrow{\mathbf{control-signal}()} \mathbf{delimit-cc}(X')} \quad (89)$$

$$\text{Rule } \mathbf{delimit-cc}(V : T) \rightsquigarrow V \quad (90)$$

The **resume-continuation** funcon takes a value as its second argument and ‘plugs’ it into the **hole** in the continuation. This is defined as follows:

$$\text{Rule } \frac{X \xrightarrow{\mathbf{plug-signal}(V)} X'}{\mathbf{resume-continuation}(\mathbf{continuation}(\mathbf{abstraction}(X)), V : T) \xrightarrow{\mathbf{plug-signal}()} X'} \quad (91)$$

$$\text{Funcon } \mathbf{hole} : \Rightarrow \mathbf{values} \quad (92)$$

$$\text{Rule } \mathbf{hole} \xrightarrow{\mathbf{plug-signal}(V)} V \quad (93)$$

5.3. Remarks

The use of **resume-continuation** and **hole** is the main simplification relative to the specification in [30]. The key enabling feature provided by CBS is control-flow entities that allow signals to synchronise between two funcons, thereby allowing information to flow in either direction. Thus **plug-signal** can transmit a value into a sub-term, and **control-signal** can transmit a value to an enclosing delimiter. In [30], output entities were used to simulate outward flowing control-flow signals (which worked rather well), and contextual entities were used to simulate inward flowing control-flow signals (which was undesirable). The main problem with the latter is that a premise cannot directly detect whether the contents of a contextual entity are used in a computation step, so it is not straightforward to enforce that a **plug-signal** is received (this is necessary in the presence of non-deterministic evaluation order and multiple occurrences of **control**). This was worked around by tagging each hole with a special kind of identifier, and then maintaining an auxiliary environment that accumulated all the values to be plugged into those holes. This was circuitous, and detracted from the clarity of the specification.

6. Executing Funcon Terms

A CBS language definition is intended to be executable so that a prototype implementation of a language can be generated from its definition. Each of the three aspects of a CBS definition – abstract syntax grammars, semantic translation functions, and the I-MSOS rules for funcons – generates a component of the prototype implementation. In Sect. 7 we discuss the tools involved in generating and connecting the components responsible for parsing a program and translating the program to a funcon term. In this section we discuss some of the features of the funcon term interpreters generated in Haskell from CBS funcon definitions by a compiler for CBS. In particular, we explain that the generated Haskell fragments are as modular as the funcon definitions that they implement. As a result, generated Haskell modules compose freely to form interpreters for arbitrary funcon collections and can be reused without recompilation. The implementation of the reusable funcon library is provided alongside CBS, enabling interpretation for all languages specified without language-specific funcons. Users employ the CBS compiler only to generate implementations for any language-specific funcons. We also discuss a suite of values built-in to CBS and the development of implementation-specific funcons and unit-tests with configuration files.

6.1. Funcon Modules

The Haskell funcon framework [31] provides three Haskell modules:

- **Funcons.EDSL** exports helper functions that are used in the implementation of funcons. For example, to modify or access semantic entities, or to test side-conditions.
- **Funcons.Tools** exports helper functions for composing *funcon modules*. A funcon module is a Haskell module exporting three collections containing information about funcon, data type, and semantic-entity definitions, respectively. Each funcon module exports an interpreter, which is aware only of the components defined in the module.
- **Funcons.Core**, a funcon module containing the implementation of the funcons in the reusable library. This includes all the funcons mentioned in this paper except the SIMPLE-specific funcon **allocate-nested-vectors**. The implementations have been generated once and for all.

Funcon modules have two noteworthy properties. Firstly, they are independent and do not need to import other funcon modules; they only need to import **Funcons.EDSL**, **Funcons.Operations** (see Sect. 6.5) and **Funcons.Tools**. Secondly, funcon modules can be freely composed: a module implementing funcons *A* and a module implementing funcons *B* are composed without restrictions to form a module implementing funcons *A* and *B*. CBS files are thus compiled individually, and the resulting modules can be composed as needed. Moreover, once a module has been compiled, it can be linked as needed, without requiring recompilation. This makes it possible to deliver the implementation of the reusable funcon library as an

```

import Funcons.Tools (mkFreshInterpreter)
import Funcons.EDSL
import Funcons.Operations hiding (Values, libFromList)
import Funcons.Tools
main = mkMainWithLibraryEntitiesTypes funcons entities types
entities = []
types = typeEnvFromList
  [("returning", DataTypeMembers "returning" []
    [DataTypeConstructor "returned"
      [TName "defined-values"] (Just [])])]
funcons = libFromList
  [("environments", NullaryFuncon stepEnvironments)
  , ("scope", NonStrictFuncon stepScope)
  , ("returning", NullaryFuncon stepReturning)
  , ("returned", StrictFuncon stepReturned)]

```

Figure 3: Excerpt from a funcon module for a contrived collection of CBS funcon specifications.

isolated package. Enforcing these properties has been a leading principle in the development of the Haskell funcon framework. For the purpose of illustration we have compiled a contrived CBS file containing the definitions of **environments**, **scope**, and **returning** (see Rules 25, 29, and 67). Figure 3 shows a fragment of the funcon module generated from this file. It omits the functions that implement the behaviour of the funcons, e.g. *stepEnvironments*.

We represent funcon terms using strings for the names of funcons in a data type, with generic constructors for funcon application, constants (funcons with no arguments), and literals (built-in values and types). A funcon module contains a *funcon library* mapping funcon names to their implementation (*funcons* in Figure 3). Similarly, we use maps to associate the names of semantic entities with values (in Figure 3, *entities* associates semantic entities with default values). Evaluating funcon terms can thus cause runtime errors because it involves searching in maps and there are no static guarantees that the required entries are available. By using strings in these ways we have gained the desired modularity, but lost static guarantees provided by a Haskell compiler. Instead, we rely on the tools explained in Sect. 7 to ensure that all mentioned funcons and entities are defined and that their code is indeed generated and included in the final interpreter.

An alternative approach would have been to use the *Data types à la carte* technique [32] to combine funcons defined in separate modules into a single data type with a constructor for each funcon. Having individual Haskell constructors for each funcon would have provided stronger correct-by-construction guarantees about the well-formedness of funcon terms. Several authors [33, 34] have used this technique for the specific purpose of defining modular programming-language constructs in a Haskell setting.

Types are either built-in to CBS (e.g. **integers** and **values**), provided as reusable components (e.g. **environments** and **identifiers**), or defined as part of a language definition. A funcon module generated from a CBS file provides a funcon implementation for each type and data-type constructor defined in the file (see *funcons* in Figure 3). In addition, a funcon module provides a type environment which relates the data types to their constructors (*types* in Figure 3). An interpreter uses the type environment for dynamic type checking.

6.2. Executable Funcons

In Sect. 4.4 we described how each class of semantic entity is implicitly propagated in CBS rules. We achieve implicit propagation in funcon code by working in a monad, and defining its *return* and *bind* operations such that they implement the desired implicit propagation. A contextual entity corresponds to a reader monad, a mutable entity corresponds to a state monad, and an output entity corresponds to a writer monad [35]. Input entities correspond to a restricted form of state monad (input is always consumed in

```

stepScope fargs =
  evalRules [rewrite1] [step1, step2]
  where
    rewrite1 = do
      let env = emptyEnv
          env ← fsMatch fargs [PAnnotated PWildcard (TName "environments")
                              , PAnnotated (PMetaVar "V") (TName "values")] env
          rewriteTermTo (TVar "V") env
    step1 = do
      let env = emptyEnv
          env ← lifted_fsMatch fargs
                [PAnnotated PWildcard (TName "environments")
                , PMetaVar "X"] env
          env ← getInhPatt "environment" (VPMetaVar "Rho0") env
          env ← lifted_sideCondition (SCPatternMatch (TApp "map-override"
                [TVar "Rho1", TVar "Rho0"]) (VPMetaVar "Rho2")) env
          env ← withInhTerm "environment" (TVar "Rho2") env
                (premise (TVar "X") (PMetaVar "X'") env)
          stepTermTo (TApp "scope" (TTuple [TVar "Rho1", TVar "X'"])) env
    step2 = do
      let env = emptyEnv
          env ← lifted_fsMatch fargs [PMetaVar "X1", PMetaVar "X2"] env
          env ← premise (TVar "X1") (PMetaVar "X1'") env
          stepTermTo (TApp "scope" (TTuple [TVar "X1'", TVar "X2"])) env

```

Figure 4: Haskell code implementing **scope**.

order, and each input can only be consumed once), and control-flow entities correspond to a combination of a reader and a writer monad (except that only one signal may be emitted in each control-flow entity at once; they do not form a monoid).

We define a data type *IMSOS* that combines the five semantic-entity classes (we omit the monad instance):

```

data IMSOS a = IMSOS (MReader → MState → (Either Exception a, MState, MWriter))
data MReader = MReader Contextual ControlFlow
data MState = MState Mutable Input
data MWriter = MWriter ControlFlow Output
data Exception = ...

```

Each semantic-entity class such as *Contextual* is implemented by a mapping from entity names to values. This avoids the need to use monad transformers [36] to add additional entities: our *IMSOS* monad is fixed. Using Haskell’s native support for monads, we can generate human-readable funcon code that only refers to the entities explicitly mentioned in the funcon definition. As a consequence, the code of a funcon is as modular and compositional as its CBS definition. Moreover, the code is only regenerated if the CBS file in which the funcon is defined is regenerated. Figure 4 gives an indication of generated funcon code. It contains the implementation of **scope**, which was omitted from Figure 3.

The code for a funcon is formed by one sequence of statements in the *IMSOS* monad for every rule of the funcon. The code of Figure 4 shows three rules. The rules *rewrite1* and *step1* correspond to Rules 31 and 30, respectively. The rule *step2* is the congruence rule generated because the signature of **scope** is strict in its first parameter. The statements can access or modify semantic-entity values, perform pattern matching and substitution, and test side conditions and premises. The statements generated for a rule may raise several kinds of internal exceptions, each representing a particular reason why the rule implemented by those statements is not applicable. For example, a funcon can be applied to incorrect arguments (e.g. if the

first argument of **if-true-else** was not a Boolean), pattern matching may have failed, or a premise may not hold. To handle failure, the *IMSOS* monad contains the *Exception* component and backtracking is used to try an alternative rule. If no rules are applicable to a particular term, i.e. all alternatives raise an exception, then the term is stuck.

The *IMSOS* monad stores extra information, including the funcon library, runtime options, and collected meta-data; but we elide those details here. The code for a funcon forms a micro-interpreter that can be executed independently, given a funcon library with entries for all funcons explicitly mentioned in the definition of the funcon. As an alternative to input/output simulation, the *IMSOS* monad also provides facilities for connecting input and output entities to real console input/output, allowing a user to run a funcon program interactively.

6.3. Context-free Rewriting

CBS supports two transition relations: computation steps and rewrites (see Sect. 4.1). Computation steps may depend on semantic entities, making them sensitive to context. Furthermore, computation steps may modify semantic entities, causing observable side effects. Conversely, the rewrite relation is context insensitive and never causes observable side effects. To separate rewriting from computation steps we use a separate monad (not shown), which does not provide access to the semantic entities. This guarantees, by construction, that rewrites do not access semantic entities.

The properties of the rewrite relation are such that, operationally, it can be applied at any time during execution anywhere in a term. Thus, funcon term interpreters are free to perform rewriting whenever is most convenient or efficient. Experience has shown that the choice of strategy to determine when rewrites are performed has a tremendous influence on the efficiency of funcon term interpreters (see also Sect. 6.7). A greedy rewriting strategy performs rewrites immediately on every occurrence of a term, possibly resulting in redundant rewrites (e.g. rewriting all arguments **if-true-else**). We choose to rewrite a term just before performing a computational step, and directly afterwards on the term resulting from the step if it was successful. This strategy exhibits inefficiencies in combination with our backtracking approach. A rule may turn out to be inapplicable after some rewrites have been performed. The rewrites are forgotten when backtracking selects the next rule to try. We expect this problem is easy to avoid in a host language providing mutable references or objects. Our Haskell implementation would benefit from graph reduction, sharing the effects of rewrites [37].

6.4. Refocussing Optimisation

Directly implementing a small-step operational semantics as the transitive closure of the computation-step relation is straightforward but inefficient. At each computation step, the interpreter traverses the funcon term from the root to a subterm (sometimes called the redex), ‘executes’ the subterm by replacing it, and reconstructs a (mostly identical) term. Simultaneously, the values held by semantic entities are maintained and possibly modified. This corresponds to constructing a derivation tree in accordance to the I-MSOS rules (including the implicit congruence rules) of a CBS language definition. The cost of each step is potentially linear in the size of the funcon term.

To overcome this inefficiency, we have added a *refocussing* optimisation [38]. The key insight of refocussing is that, in many situations, once execution of a subterm has begun, the subsequent computation steps will involve executing that same subterm, until it is reduced to a value. Therefore, rather than re-traversing the term each step, the interpreter can ‘take a short-cut’ and locally continue executing the subterm. This optimisation is easily implemented and significantly improves the efficiency of the interpreter.

However, it is not always valid to ‘take a short-cut’. In general, deciding the next subterm to be executed requires checking the conditions of the inference rules. The execution of a computational step, which involves replacing some subterm r with r' , may be accompanied by a modification to the value held by some semantic entity. Because of the latter, the conditions for choosing r as the subterm to be executed may no longer hold, in which case the next subterm to be executed is not r' , invalidating refocussing. For example, when a control-flow signal is emitted, execution of the current subterm should *not* continue, as that signal needs to be handled further up the term. Typically, the handler for a control-flow signal will then change the

subterm that is next to be executed (e.g. Rule 63 for **handle-thrown** in Sect. 4.4.5). Our interpreter checks for emitted signals and additional output to determine whether to apply refocussing. To guarantee the validity of refocussing in all cases, we place restrictions on the way semantic entities are used in inference rules, as suggested by Bach Poulsen [39].

6.5. Values and Value Operations

The CBS framework provides a large suite of values and operations on those values (see Appendix B.2 for examples). Some operations have not been specified in CBS, for example operations on integers and floating-point numbers. In these cases the operations are manually implemented and exported by the module `Funcons.Operations` provided by a separate library [40]. The library is separate from the the Haskell funcon framework and can be used by other semantic frameworks implemented in Haskell.

Some built-in values are native to Haskell, such as floating-point numbers and ASCII characters. Other types of values are provided by libraries, such as vectors and multisets.¹⁰ Most of the value operations are directly implemented by a function provided by an existing Haskell library. In other cases, the operation is easy to define in terms of existing operations or functions. Creating funcon modules that export the corresponding funcons is done manually by unwrapping values from the funcon term data type, applying the corresponding operation from `Funcons.Operations`, and wrapping the result. This work is done once and for all, as the suite of values and operations is fixed (exported by `Funcons.Core`).

6.6. Configuration

The interpreters exported by funcon modules are easily configured by command line flags or configuration files. For example, a user can restrict the number of computational steps performed. This is useful to determine which steps lead up to unwanted behaviour (arising from an invalid program or an incorrect specification). The interpreters can print debugging information such as the values of semantic entities, the number of computational steps performed, or the number of rewrites performed. The documentation of `Funcons.Tools` [31] contains an exhaustive enumeration of all configuration options.

Configuration files offer a limited method for defining funcons. A particular use case is to define *implementation-specific* funcons representing a particular implementation choice. Language manuals often leave options for language implementers regarding certain aspects of the language, e.g., the number of bits used in the binary representation of different forms of numbers. The following excerpt shows the definition of some implementation-specific funcons in a configuration file:

```
funcons {
  implemented-floats-format = binary64;
  implemented-integers-width = 31;
  implemented-characters    = unicode-characters;
}
```

Configuration files also offer a simple but effective method for defining unit tests. A configuration file for a unit test contains an initial term to be evaluated and the expected outcome of evaluating the term. The expected outcome consists of a result term, as well as the values of some semantic entities. The excerpt below shows an example configuration file containing a unit test.

```
general{
  funcon-term: give(read,sequential(print(given),print(given)));
}
inputs{
  standard-in: 2,3; // provides 2 and 3 as input
}
tests{
  result-term: null; // expects null as the result value
  standard-out: 2,2; // expects 2 is printed twice
  standard-in: 3; // expects 3 as remaining input
}
```

¹⁰The vector and multiset packages on Hackage.

Test Program	Generated Interpreter		K Tool
	Unoptimised	Refocussing Enabled	
exception tests 1 to 15	11.4	1.67	30.3
div-nondet	0.2	0.06	1.9
factorial	2.8	0.18	1.9
collatz	11.3	0.53	1.9
running total	25.6	2.43	35.9
higher-order	-	3.02	10.1
matrix	-	5.84	2.2
sortings	-	6.11	2.9
running total	-	17.4	51.1

Table 1: Runtimes in seconds of the generated SIMPLE interpreter on several test programs.

Differences between expected and actual outcome are reported to the standard output. A successful unit test is silent. Entities not mentioned in a unit test are treated differently depending on the class of the entity. For example, any value for **store** is considered acceptable outcome, if the tests does not mention **store**. However, a raised **abrupted** signal is considered a violation if the test did not mention the entity **abrupted**.

6.7. Evaluation

The Haskell funcon framework is developed together with a CBS compiler that generates funcon modules from CBS specifications. Using these tools, an interpreter is obtained in two steps. Firstly, the CBS compiler is used to generate code for any language-specific funcons defined in the specification. Secondly, using helper functions in `Funcons.Tools` of the framework, the funcon modules generated for the language are combined with `Funcons.Core`, the module containing the implementations of the reusable funcons. In the case of SIMPLE, the resulting funcon module exports an interpreter capable of executing all funcon terms produced by translating SIMPLE programs. The interpreter for SIMPLE is available on Hackage [41]. If a specification does not define language-specific funcons, then no code needs to be generated and the interpreter provided by `Funcons.Core` can be used directly.

We have tested the SIMPLE interpreter on some of the SIMPLE test programs that come bundled with distributions of the K tool [42]. Table 1 presents the runtimes of the interpreter for several of these test programs, with and without the refocussing optimisation enabled (see Sect. 6.4). For reference, we have included the runtimes of the same test programs executed with Version 5 of the K tool. These results were produced on a virtual private server with four 2.2GHz virtual CPUs and 8GiB of RAM, under Ubuntu 18.04. The Haskell funcon modules were compiled using GHC 8.0.2. The K specification of SIMPLE was compiled with the Java back end of the K tool and loaded in server-mode to avoid compilation overhead on each run. Several runs of ‘JVM warm up’ have been performed. The numbers are the average runtimes of 10 runs.

Without refocussing enabled, the tests **higher-order**, **matrix** and **sortings** cause the memory to overflow. These tests involve a relatively large number of function calls or loop-iterations, which result in considerable growth of the funcon term under evaluation, thus increasing the overhead of decomposing and recomposing the term at each step of the small-step evaluation. Decomposition involves potentially a lot of backtracking and undoing context-free rewrites unnecessarily, as discussed in Sect. 6.3. With refocussing, these programs execute in three to six seconds.

This experiment shows that the generated SIMPLE interpreter is adequate for testing the SIMPLE specification by running basic programs.

7. Integrated Development Environment for CBS

CBS language definitions can be developed, maintained, and tested in Eclipse. We have used SPOOFAX [8] to develop an Eclipse-based editor for CBS language definitions, with many useful features for syntax-aware browsing and editing.

- The CBS editor has menu actions for generating language editors, web pages, and funcon indexes from CBS language definitions; outdated files are detected and regenerated when changes have been made.
- A generated language editor parses programs in the specified language (flagging any text that has syntax errors or ambiguities) and provides basic syntax highlighting; there is also a menu action for translating programs into funcon terms.
- The CBS compiler generates modules for any language-specific funcons from CBS language definitions. These are composed with `Funcons.Core` to form a language-specific funcon term interpreter. The CBS compiler and the interpreter are added as external tools so that they can be used from within Eclipse.

The rest of this section discusses these features in more detail.

7.1. Editing CBS Files

The CBS editor analyses all CBS source files in an Eclipse project, updating the analysis in the background whenever the user edits the text. It checks that the text conforms to the context-free grammar of the CBS meta-language, adds syntax highlighting, and flags as erroneous any parts of the text that cannot be parsed. It also creates hyperlinks from uses of symbols to the locations where the symbols are introduced, flagging missing symbols as errors.

The CBS analysis is independent of the division of language and funcon definitions into files and folders. For example, we may divide language definitions into sections in the same way as reference manuals, and put each section in a separate file. The reusable funcons are currently grouped together in files according to the semantic entities involved, to facilitate browsing of closely related funcons. The hyperlinks in the edited files allow navigation and browsing of CBS specifications without awareness of the underlying file system. The analysis was implemented by specifying it in the current SPOOFAX meta-language for static semantics, NaBL2 [43], based on the concepts of scope graphs [44].

When a programming language evolves, the syntax and/or semantics of its constructs can change, new constructs may be added, and existing constructs may be removed. This is achieved by editing the CBS files that define the language syntax and its translation to funcons. If a new funcon is needed, it can be added to the language specification for local use (like **allocate-nested-vectors**). New funcons that are language-independent may subsequently be submitted for inclusion in the library of reusable funcons.

The screenshot in Figure 5 shows several files open in our integrated development environment (IDE) during the development of the component-based semantics of SIMPLE. The top-left pane is browsing the abstract syntax and semantics of variable declarations. In the bottom-left pane, a CBS rule defining the translation of variable-declaration statements to funcons is being edited; the red mark in the margin flags an error in the name of the translation function for declarations. The colours and fonts distinguish the names of syntax non-terminals (*green*), funcons (*red*), semantic functions (*blue italic*), and variables (*black italic*). Clicking on a name in a CBS editor shows its definition in a separate pane. The top-right pane shows the CBS file containing the definition of the reusable funcon **scope**.

Name analysis and arity checking were specified in NaBL2: the editor resolves all references to names (with separate namespaces for funcons, syntax sorts, translation functions, and meta-variables) to their declarations, flagging any missing or duplicate declarations as errors. However, full type checking for CBS cannot be specified in NaBL2, which does not support structural subtyping.

7.2. Generating Funcon Terms

A generated language editor uses a scannerless generalised LR (SGLR) parser generated from the abstract syntax grammar of a CBS language definition. An abstract syntax grammar is usually highly ambiguous; to support disambiguation, disambiguation rules are currently written inside comment blocks in CBS specifications. The language editor has a menu action for translating the abstract syntax trees of disambiguated programs to funcon terms, executing code generated from the semantic equations. The two panes on the lower right of Figure 5 show a small test program and part of its translation to funcons.




Figure 5: The IDE for CBS in action.

7.3. Evaluating Funcon Terms

Funcon term interpreters can be added to Eclipse as ‘external tools’. This makes it possible to execute funcon term files and configuration files with unit tests from within Eclipse, showing the output of the interpreter in Eclipse’s console. If a language specification defines any language-specific funcons, then the CBS files that define these funcons need to be compiled with the CBS compiler to produce interpreter code. The CBS compiler can be added to Eclipse as an external tool as well. The CBS compiler generates modular interpreters, as described in Sect. 6. If a CBS file containing a language-specific funcon is modified, only that file needs to be recompiled. The CBS compiler is not needed for any projects that use only funcons from the reusable library.

8. Conclusion

We have presented CBS, a unified meta-language for the formal specification of programming languages. Our intent is that a language designer will use CBS to define the syntax and semantics of a source language, using context-free grammars and translation functions that produce funcon terms (Sect. 3). The funcons themselves are defined in CBS and are provided as part of a library of over 100 pre-defined funcons that can be reused in the development of different programming languages. The reuse of formally specified components is the key novelty compared to other specification frameworks. If required, language-specific funcons can be defined, and new reusable funcons can be added to the main (open-ended) funcon library; this is possible because of the *modular* nature of the underlying formal framework (MSOS).

CBS definitions are *executable*. This allows the definitions of funcons and programming languages to be validated through testing, and for prototype interpreters to be generated from CBS definitions of new languages. We provide a suite of tools to support interactive development of funcon and language definitions (Sect. 7) as well as generating funcon term interpreters (Sect. 6).

Our hope is that with the release of over 100 validated funcons, supported by several case studies, CBS could have significant transformative effects: language developers could be encouraged to make use of formal semantics and experiment with generated prototype implementations during the design process; domain experts could be empowered to design, specify and implement domain-specific languages themselves; and researchers and students would be provided with a dedicated portal and online open-access repository for language and funcon definitions.

Acknowledgements

The authors are grateful to the anonymous reviewers for helpful comments and suggestions for improvement. The reported work was supported by EPSRC grants for the PLANCOMPS project at Swansea University (EP/I032495/1) and Royal Holloway, University of London (EP/I032509/1).

References

- [1] R. Milner, M. Tofte, R. Harper, D. MacQueen, *The Definition of Standard ML (Revised)*, MIT Press, Cambridge, MA, USA, 1997.
- [2] P. Hudak, J. Hughes, S. L. Peyton Jones, P. Wadler, A history of Haskell: Being lazy with class, in: *Third Conference on History of Programming Languages*, ACM, 2007, pp. 12:1–12:55. doi:10.1145/1238844.1238856.
- [3] P. D. Mosses, M. J. New, Implicit propagation in structural operational semantics, in: *Fifth Workshop on Structural Operational Semantics*, Vol. 229(4) of *Electronic Notes in Theoretical Computer Science*, Elsevier, 2009, pp. 49–66. doi:10.1016/j.entcs.2009.07.073.
- [4] P. D. Mosses, Modular structural operational semantics, *Journal of Logic and Algebraic Programming* 60–61 (2004) 195–228. doi:10.1016/j.jlap.2004.03.008.
- [5] M. Churchill, P. D. Mosses, N. Sculthorpe, P. Torrini, Reusable components of semantic specifications, in: *Transactions on Aspect-Oriented Software Development XII*, Vol. 8989 of *Lecture Notes in Computer Science*, Springer, 2015, pp. 132–179. doi:10.1007/978-3-662-46734-3_4.
- [6] X. Leroy, *Caml Light manual*, <http://caml.inria.fr/pub/docs/manual-caml-light> (1997).
- [7] PLANCompS project, CBS and funcons beta release, GitHub, [Online, accessed 17th December 2018] (2018). URL <https://plancomps.github.io/CBS-beta>

- [8] L. C. L. Kats, E. Visser, The Spoofax language workbench: Rules for declarative specification of languages and IDEs, in: International Conference on Object Oriented Programming Systems Languages and Applications, ACM, 2010, pp. 444–463. doi:10.1145/1869459.1869497.
- [9] G. Roşu, T. F. Şerbănuţă, \mathbb{K} overview and SIMPLE case study, *Electronic Notes in Theoretical Computer Science* 304 (2014) 3–56. doi:10.1016/j.entcs.2014.05.002.
- [10] G. D. Plotkin, A structural approach to operational semantics, *Journal of Logic and Algebraic Programming* 60–61 (2004) 17–139, reprint of Technical Report FN-19, DAIMI, Aarhus University, 1981. doi:10.1016/j.jlap.2004.05.001.
- [11] M. Churchill, P. D. Mosses, Modular bisimulation theory for computations and values, in: 16th International Conference on Foundations of Software Science and Computation Structures, Vol. 7794 of Lecture Notes in Computer Science, Springer, 2013, pp. 97–112. doi:10.1007/978-3-642-37075-5_7.
- [12] P. D. Mosses, F. Vesely, Weak bisimulation as a congruence in MSOS, in: Logic, Rewriting, and Concurrency, Vol. 9200 of Lecture Notes in Computer Science, Springer, 2015, pp. 519–538. doi:10.1007/978-3-319-23165-5_24.
- [13] F. Chalub, C. Braga, Maude MSOS tool, *Electronic Notes in Theoretical Computer Science* 176 (4) (2007) 133–146. doi:10.1016/j.entcs.2007.06.012.
- [14] P. D. Mosses, Teaching semantics of programming languages with Modular SOS, in: Teaching Formal Methods 2006: Practice and Experience, Electronic Workshops in Computing, BCS, 2006, pp. 275–286. URL <https://ewic.bcs.org/content/ConWebDoc/9093>
- [15] L. T. van Binsbergen, N. Sculthorpe, P. D. Mosses, Tool support for component-based semantics, in: Companion Proceedings of the 15th International Conference on Modularity, ACM, 2016, pp. 8–11. doi:10.1145/2892664.2893464.
- [16] R. Farahbod, V. Gervasi, U. Glässer, CoreASM: An extensible ASM execution engine, *Fundamenta Informaticae* 77 (1–2) (2007) 71–103.
- [17] P. Sewell, F. Z. Nardelli, S. Owens, G. Peskine, T. Ridge, S. Sarkar, R. Strniša, Ott: Effective tool support for the working semanticist, *Journal of Functional Programming* 20 (1) (2010) 71–122. doi:10.1017/S0956796809990293.
- [18] M. Felleisen, R. B. Findler, M. Flatt, *Semantics Engineering with PLT Redex*, MIT Press, 2009.
- [19] G. Roşu, T. F. Şerbănuţă, An overview of the \mathbb{K} semantic framework, *Journal of Logic and Algebraic Programming* 79 (6) (2010) 397–434. doi:10.1016/j.jlap.2010.03.012.
- [20] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, J. F. Quesada, Maude: specification and programming in rewriting logic, *Theoretical Computer Science* 285 (2) (2002) 187–243. doi:10.1016/S0304-3975(01)00359-0.
- [21] T. Degueule, B. Combemale, A. Blouin, O. Barais, J.-M. Jézéquel, Melange: A meta-language for modular and reusable development of DSLs, in: Eighth International Conference on Software Language Engineering, ACM, 2015, pp. 25–36. doi:10.1145/2814251.2814252.
- [22] V. Vergu, P. Neron, E. Visser, DynSem: A DSL for dynamic semantics specification, in: 26th International Conference on Rewriting Techniques and Applications, Vol. 36 of Leibniz International Proceedings in Informatics, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015, pp. 365–378. doi:10.4230/LIPIcs.RTA.2015.365.
- [23] P. D. Mosses, F. Vesely, FunKons: Component-based semantics in \mathbb{K} , in: 10th International Workshop on Rewriting Logic and Its Applications, Vol. 8663 of Lecture Notes in Computer Science, Springer, 2014, pp. 213–229. doi:10.1007/978-3-319-12904-4_12.
- [24] P. D. Mosses, A modular SOS for ML concurrency primitives, BRICS Research Series RS-99-57, Department of Computer Science, Aarhus University (1999). URL <http://www.brics.dk/RS/99/57/>
- [25] D. W. Barron, J. N. Buxton, D. F. Hartley, E. Nixon, C. Strachey, The main features of CPL, *The Computer Journal* 6 (2) (1963) 134–143. doi:10.1093/comjnl/6.2.134.
- [26] C. Strachey, Fundamental concepts in programming languages, *Higher-Order and Symbolic Computation* 13 (1) (2000) 11–49. doi:10.1023/A:1010000313106.
- [27] P. D. Mosses, Pragmatics of modular SOS, in: International Conference on Algebraic Methodology and Software Technology, Vol. 2422 of Lecture Notes in Computer Science, Springer, 2002, pp. 21–40. doi:10.1007/3-540-45719-4_3.
- [28] O. Danvy, A. Filinski, A functional abstraction of typed contexts, Tech. Rep. 89/12, DIKU, University of Copenhagen (1989). URL <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.43.4822>
- [29] M. Felleisen, The theory and practice of first-class prompts, in: 15th Symposium on Principles of Programming Languages, ACM, 1988, pp. 180–190. doi:10.1145/73560.73576.
- [30] N. Sculthorpe, P. Torrini, P. D. Mosses, A modular structural operational semantics for delimited continuations, in: Post-proceedings of the 2015 Workshop on Continuations, Vol. 212 of Electronic Proceedings in Theoretical Computer Science, Open Publishing Association, 2016, pp. 63–80. doi:10.4204/EPTCS.212.5.
- [31] L. T. van Binsbergen, N. Sculthorpe, The Haskell funcon framework, Hackage, [Online, accessed 17th December 2018] (2018). URL <https://hackage.haskell.org/package/funcons-tools-0.2.0.5>
- [32] W. Swierstra, Data types à la carte, *Journal of Functional Programming* 18 (4) (2008) 423–436. doi:10.1017/S0956796808006758.
- [33] L. E. Day, G. Hutton, Towards modular compilers for effects, in: 12th International Symposium on Trends in Functional Programming, Vol. 7193 of Lecture Notes in Computer Science, Springer, 2012, pp. 49–64. doi:10.1007/978-3-642-32037-8_4.
- [34] N. Wu, T. Schrijvers, R. Hinze, Effect handlers in scope, in: 2014 Symposium on Haskell, ACM, 2014, pp. 1–12. doi:10.1145/2633357.2633358.
- [35] M. P. Jones, Functional programming with overloading and higher-order polymorphism, in: First International School

- on Advanced Functional Programming, Vol. 925 of Lecture Notes in Computer Science, Springer, 1995, pp. 97–136. doi:10.1007/3-540-59451-5_4.
- [36] S. Liang, P. Hudak, M. Jones, Monad transformers and modular interpreters, in: 22nd Symposium on Principles of Programming Languages, ACM, 1995, pp. 333–343. doi:10.1145/199448.199528.
- [37] S. L. Peyton Jones, The Implementation of Functional Programming Languages, Prentice Hall, 1987.
- [38] O. Danvy, L. R. Nielsen, Refocusing in reduction semantics, BRICS Research Series RS-04-26, Department of Computer Science, Aarhus University (2004).
URL <http://www.brics.dk/RS/04/26/>
- [39] C. Bach Poulsen, Extensible transition system semantics, Ph.D. thesis, Swansea University (2016).
- [40] L. T. van Binsbergen, N. Sculthorpe, Funcons values and value operations, Hackage, [Online, accessed 17th December 2018] (2018).
URL <https://hackage.haskell.org/package/funcons-values-0.1.0.3>
- [41] L. T. van Binsbergen, N. Sculthorpe, Funcon interpreter for the SIMPLE language, Hackage, [Online, accessed 17th December 2018] (2018).
URL <https://hackage.haskell.org/package/funcons-simple-0.1.0.3>
- [42] The \mathbb{K} framework, GitHub, [Online, accessed 17th December 2018] (2018).
URL <https://github.com/kframework/k5>
- [43] H. van Antwerpen, P. Néron, A. P. Tolmach, E. Visser, G. Wachsmuth, A constraint language for static semantic analysis based on scope graphs, in: 2016 Workshop on Partial Evaluation and Program Manipulation, ACM, 2016, pp. 49–60. doi:10.1145/2847538.2847543.
- [44] P. Néron, A. P. Tolmach, E. Visser, G. Wachsmuth, A theory of name resolution, in: 24th European Symposium on Programming Languages and Systems, Vol. 9032 of Lecture Notes in Computer Science, Springer, 2015, pp. 205–231. doi:10.1007/978-3-662-46669-8_9.

Appendix A. Simple Language Definition

This appendix presents our CBS definition of the SIMPLE language. The top-level translation function is *run*. We omit the lexical grammar for the non-terminals *bool*, *int*, *string* and *id*, and the equations for corresponding translation functions *val* and *id*, but they are available to be viewed online [7].

Appendix A.1. Values

Syntax $V : \text{value} ::= \text{bool} \mid \text{int} \mid \text{string}$

Semantics $\text{val}[_ : \text{value}] : \Rightarrow \text{values}$

Semantics $\text{id}[_ : \text{id}] : \Rightarrow \text{ids}$

Appendix A.2. Expressions

Syntax $\text{Exp} : \text{exp} ::= \text{'(' exp ')'} \mid \text{value} \mid \text{lexp} \mid \text{lexp '=' exp} \mid \text{'++' lexp}$

$\mid \text{'-' exp} \mid \text{exp '(' exps? ')'} \mid \text{'sizeOf' '(' exp ')'} \mid \text{'read' '(' ')'} \mid$

$\mid \text{exp '+' exp} \mid \text{exp '-' exp} \mid \text{exp '*' exp} \mid \text{exp '/' exp} \mid \text{exp '%' exp}$

$\mid \text{exp '<' exp} \mid \text{exp '<=' exp} \mid \text{exp '>' exp} \mid \text{exp '>=' exp}$

$\mid \text{exp '==' exp} \mid \text{exp '!=' exp} \mid \text{'!' exp} \mid \text{exp '&\&' exp} \mid \text{exp '||' exp}$

Rule $\llbracket \text{'(' Exp ')'} \rrbracket : \text{exp} = \llbracket \text{Exp} \rrbracket$

Semantics $rval[- : exp] : \Rightarrow \text{values}$

Rule $rval[V] = val[V]$

Rule $rval[LExp] = \text{assigned}(lval[LExp])$

Rule $rval[LExp '=' Exp] =$

give($rval[Exp]$, **sequential**(**assign**($lval[LExp]$, **given**), **given**))

Rule $rval['++' LExp] =$

give($lval[LExp]$, **sequential**(**assign**(**given**, **integer-add**(**assigned**(**given**), 1)), **assigned**(**given**)))

Rule $rval['-' Exp] = \text{integer-negate}(rval[Exp])$

Rule $rval[Exp '(' Exps ')'] = \text{apply}(rval[Exp], \text{tuple}(rvals[Exps]))$

Rule $rval['\text{sizeOf}' '(' Exp ')'] = \text{length}(\text{vector-elements}(rval[Exp]))$

Rule $rval['\text{read}' '(' ')'] = \text{read}$

Rule $rval[Exp_1 '+' Exp_2] = \text{integer-add}(rval[Exp_1], rval[Exp_2])$

Rule $rval[Exp_1 '-' Exp_2] = \text{integer-subtract}(rval[Exp_1], rval[Exp_2])$

Rule $rval[Exp_1 '*' Exp_2] = \text{integer-multiply}(rval[Exp_1], rval[Exp_2])$

Rule $rval[Exp_1 '/' Exp_2] = \text{checked}(\text{integer-divide}(rval[Exp_1], rval[Exp_2]))$

Rule $rval[Exp_1 \% Exp_2] = \text{checked}(\text{integer-modulo}(rval[Exp_1], rval[Exp_2]))$

Rule $rval[Exp_1 '<' Exp_2] = \text{is-less}(rval[Exp_1], rval[Exp_2])$

Rule $rval[Exp_1 '<=' Exp_2] = \text{is-less-or-equal}(rval[Exp_1], rval[Exp_2])$

Rule $rval[Exp_1 '>' Exp_2] = \text{is-greater}(rval[Exp_1], rval[Exp_2])$

Rule $rval[Exp_1 '>=' Exp_2] = \text{is-greater-or-equal}(rval[Exp_1], rval[Exp_2])$

Rule $rval[Exp_1 '==' Exp_2] = \text{is-equal}(rval[Exp_1], rval[Exp_2])$

Rule $rval[Exp_1 '!=' Exp_2] = \text{not}(\text{is-equal}(rval[Exp_1], rval[Exp_2]))$

Rule $rval['!' Exp] = \text{not}(rval[Exp])$

Rule $rval[Exp_1 '&\&' Exp_2] = \text{if-true-else}(rval[Exp_1], rval[Exp_2], \text{false})$

Rule $rval[Exp_1 '||' Exp_2] = \text{if-true-else}(rval[Exp_1], \text{true}, rval[Exp_2])$

Syntax $Exps : \text{exps} ::= \text{exp} (' , ' \text{exps})^?$

Semantics $rvals[- : \text{exps}^?] : (\Rightarrow \text{values})^*$

Rule $rvals[] = ()$

Rule $rvals[Exp] = rval[Exp]$

Rule $rvals[Exp ', Exps] = rval[Exp], rvals[Exps]$

Syntax $LExp : \text{lexp} ::= \text{id} | \text{lexp} '[' \text{exps} ']'$

Rule $[[LExp '[' Exp ', Exps ']] : \text{lexp} = [[LExp '[' Exp ']' '[' Exps ']]$

Semantics $lval[- : \text{lexp}] : \Rightarrow \text{variables}$

Rule $lval[Id] = \text{bound}(id[Id])$

Rule $lval[LExp '[' Exp ']] = \text{checked}(\text{index}(\text{integer-add}(1, rval[Exp]), \text{vector-elements}(rval[LExp])))$

Appendix A.3. Statements

Syntax Block : block ::= '{' stmts[?] '}'

Syntax Stmts : stmts ::= stmt stmts[?]

Syntax Stmt : stmt ::= imp-stmt | vars-decl

Syntax ImpStmt : imp-stmt ::= block | exp ';' |

| 'if' '(' exp ')' block ('else' block)[?] | 'while' '(' exp ')' block
 | 'for' '(' stmt exp ';' exp ')' block | 'print' '(' exps ')' ';' | 'return' exp[?] ';' |
 | 'try' block 'catch' '(' id ')' block | 'throw' exp ';' |

Rule $\llbracket \text{'if' '(' Exp ')' Block} \rrbracket$: stmt = $\llbracket \text{'if' '(' Exp ')' Block 'else' '{' '}' } \rrbracket$

Rule $\llbracket \text{'for' '(' Stmt Exp_1 ';' Exp_2 ')' '{' Stmts '}' } \rrbracket$: stmt =
 $\llbracket \text{'{' Stmt 'while' '(' Exp_1 ')' '{' '{' Stmts '}' Exp_2 ';' '}' '}' } \rrbracket$

Semantics exec $\llbracket _ : stmts \rrbracket$: \Rightarrow **null-type**

Rule $\text{exec} \llbracket \text{' ' } \rrbracket = \mathbf{null}$

Rule $\text{exec} \llbracket \text{'{' Stmts '}' } \rrbracket = \text{exec} \llbracket \text{Stmts} \rrbracket$

Rule $\text{exec} \llbracket \text{ImpStmt Stmts} \rrbracket = \mathbf{sequential}(\text{exec} \llbracket \text{ImpStmt} \rrbracket, \text{exec} \llbracket \text{Stmts} \rrbracket)$

Rule $\text{exec} \llbracket \text{VarsDecl Stmts} \rrbracket = \mathbf{scope}(\text{declare} \llbracket \text{VarsDecl} \rrbracket, \text{exec} \llbracket \text{Stmts} \rrbracket)$

Rule $\text{exec} \llbracket \text{VarsDecl} \rrbracket = \mathbf{effect}(\text{declare} \llbracket \text{VarsDecl} \rrbracket)$

Rule $\text{exec} \llbracket \text{Exp ';' } \rrbracket = \mathbf{effect}(\text{rval} \llbracket \text{Exp} \rrbracket)$

Rule $\text{exec} \llbracket \text{'if' '(' Exp ')' Block_1 'else' Block_2} \rrbracket =$

if-true-else($\text{rval} \llbracket \text{Exp} \rrbracket, \text{exec} \llbracket \text{Block}_1 \rrbracket, \text{exec} \llbracket \text{Block}_2 \rrbracket$)

Rule $\text{exec} \llbracket \text{'while' '(' Exp ')' Block} \rrbracket = \mathbf{while}(\text{rval} \llbracket \text{Exp} \rrbracket, \text{exec} \llbracket \text{Block} \rrbracket)$

Rule $\text{exec} \llbracket \text{'print' '(' Exps ')' ';' } \rrbracket = \mathbf{print}(\text{rvals} \llbracket \text{Exps} \rrbracket)$

Rule $\text{exec} \llbracket \text{'return' Exp ';' } \rrbracket = \mathbf{return}(\text{rval} \llbracket \text{Exp} \rrbracket)$

Rule $\text{exec} \llbracket \text{'return' ';' } \rrbracket = \mathbf{return}(\mathbf{null})$

Rule $\text{exec} \llbracket \text{'try' Block_1 'catch' '(' Id ')' Block_2} \rrbracket =$

handle-thrown($\text{exec} \llbracket \text{Block}_1 \rrbracket,$

scope($\text{bind}(\text{id} \llbracket \text{Id} \rrbracket, \text{allocate-initialised-variable}(\text{values, given})), \text{exec} \llbracket \text{Block}_2 \rrbracket$))

Rule $\text{exec} \llbracket \text{'throw' Exp ';' } \rrbracket = \mathbf{throw}(\text{rval} \llbracket \text{Exp} \rrbracket)$

Appendix A.4. Declarations

Syntax Decl : decl ::= vars-decl | func-decl

Semantics declare $\llbracket _ : decl \rrbracket$: \Rightarrow **environments**

Appendix A.4.1. Variable Declarations

Syntax VarsDecl : vars-decl ::= 'var' declarators ';' |

Syntax Declarators : declarators ::= declarator (';' declarators)[?]

Rule $\llbracket \text{'var' Declarator ';' Declarators ';' Stmts} \rrbracket$: stmts =

$\llbracket \text{'var' Declarator ';' 'var' Declarators ';' Stmts} \rrbracket$

Rule $\llbracket \text{'var' Declarator ';' Declarators ';' Decl} \rrbracket$: decls =

$\llbracket \text{'var' Declarator ';' 'var' Declarators ';' Decl} \rrbracket$

Rule $\text{declare} \llbracket \text{'var' Declarator ';' } \rrbracket = \text{var-declare} \llbracket \text{Declarator} \rrbracket$

Syntax Declarator : declarator ::= id | id '=' exp | id ranks

Semantics var-declare $\llbracket _ : \text{declarator} \rrbracket : \Rightarrow \text{environments}$

Rule *var-declare* $\llbracket Id \rrbracket = \text{bind}(id \llbracket Id \rrbracket, \text{allocate-variable}(\text{values}))$

Rule *var-declare* $\llbracket Id '=' Exp \rrbracket =$

$\text{bind}(id \llbracket Id \rrbracket, \text{allocate-initialised-variable}(\text{values}, \text{rval} \llbracket Exp \rrbracket))$

Rule *var-declare* $\llbracket Id \text{ Ranks} \rrbracket = \text{bind}(id \llbracket Id \rrbracket, \text{allocate-nested-vectors}(\text{ranks} \llbracket Ranks \rrbracket))$

Appendix A.4.2. Array Declarations

Syntax Ranks : ranks ::= '[' exps ']' ranks[?]

Rule $\llbracket '[' Exp ', Exps ']' Ranks^? \rrbracket : \text{ranks} = \llbracket '[' Exp ']' '[' Exps ']' Ranks^? \rrbracket$

Semantics ranks $\llbracket _ : \text{ranks} \rrbracket : (\Rightarrow \text{nats})^+$

Rule *ranks* $\llbracket '[' Exp ']' \rrbracket = \text{rval} \llbracket Exp \rrbracket$

Rule *ranks* $\llbracket '[' Exp ']' Ranks \rrbracket = \text{rval} \llbracket Exp \rrbracket, \text{ranks} \llbracket Ranks \rrbracket$

Funcon allocate-nested-vectors $(_ : \text{nats}^+) : \Rightarrow \text{variables}$

Rule *allocate-nested-vectors* $(N : \text{nats}) \rightsquigarrow$

$\text{allocate-initialised-variable}(\text{vectors}(\text{variables}),$
 $\text{vector}(\text{left-to-right-repeat}(\text{allocate-variable}(\text{values}), 1, N)))$

Rule *allocate-nested-vectors* $(N : \text{nats}, N^+ : \text{nats}^+) \rightsquigarrow$

$\text{allocate-initialised-variable}(\text{vectors}(\text{variables}),$
 $\text{vector}(\text{left-to-right-repeat}(\text{allocate-nested-vectors}(N^+), 1, N)))$

Appendix A.4.3. Function Declarations

Syntax FuncDecl : func-decl ::= 'function' id '(' ids[?] ')' block

Rule *declare* $\llbracket \text{'function' } Id \text{ '(' } Ids^? \text{ ')' } Block \rrbracket =$

$\text{bind}(id \llbracket Id \rrbracket, \text{allocate-variable}(\text{functions}(\text{tuples}(\text{values}^*), \text{values})))$

Semantics initialise $\llbracket _ : \text{decl} \rrbracket : \Rightarrow \text{null-type}$

Rule *initialise* $\llbracket \text{'var' } Declarators \text{ ';' } \rrbracket = \text{null}$

Rule *initialise* $\llbracket \text{'function' } Id \text{ '(' } Ids^? \text{ ')' } Block \rrbracket =$

$\text{assign}(\text{bound}(id \llbracket Id \rrbracket), \text{function}(\text{closure}(\text{scope}(\text{match}(\text{given}, \text{tuple}(\text{patts} \llbracket Ids^? \rrbracket)), \text{handle-return}(\text{exec} \llbracket Block \rrbracket))))))$

Syntax Ids : ids ::= id (' , ' ids)[?]

Semantics patts $\llbracket _ : \text{ids}^? \rrbracket : \text{patterns}^*$

Rule *patts* $\llbracket \rrbracket = ()$

Rule *patts* $\llbracket Id \rrbracket = \text{pattern}(\text{closure}(\text{bind}(id \llbracket Id \rrbracket, \text{allocate-initialised-variable}(\text{values}, \text{given}))))$

Rule *patts* $\llbracket Id \text{ ';' } Ids \rrbracket = \text{patts} \llbracket Id \rrbracket, \text{patts} \llbracket Ids \rrbracket$

Appendix A.5. Programs

Syntax $Pgm : pgm ::= decls$

Semantics $run [- : pgm] : \Rightarrow \mathbf{values}$

Rule $run [Decls] = \mathbf{scope}(\mathbf{collateral}(declarations [Decls]),$
 $\mathbf{sequential}(initialisations [Decls], \mathbf{apply}(\mathbf{assigned}(\mathbf{bound}(\mathbf{"main"})), \mathbf{tuple}()))))$

Syntax $Decls : decls ::= decl decls^?$

Semantics $declarations [- : decls] : (\Rightarrow \mathbf{environments})^+$

Rule $declarations [Decl] = \mathbf{declare} [Decl]$

Rule $declarations [Decl Decls] = \mathbf{declare} [Decl], declarations [Decls]$

Semantics $initialisations [- : decls] : (\Rightarrow \mathbf{null-type})^+$

Rule $initialisations [Decl] = \mathbf{initialise} [Decl]$

Rule $initialisations [Decl Decls] = \mathbf{initialise} [Decl], initialisations [Decls]$

Appendix B. Funcon Signatures

This appendix lists the signatures of all funcons and value operations that appear in this article.

Appendix B.1. Funcons

abrupt($_ : \text{values}$) : $\Rightarrow \text{empty-type}$
allocate-initialised-variable($T : \text{types}, _ : T$) : $\Rightarrow \text{variables}$
allocate-nested-vectors($_ : \text{nats}^+$) : $\Rightarrow \text{variables}$
allocate-variable($_ : \text{types}$) : $\Rightarrow \text{variables}$
apply($_ : \text{functions}(S, T), _ : S$) : $\Rightarrow T$
assign($_ : \text{variables}, _ : \text{values}$) : $\Rightarrow \text{null-type}$
assigned($_ : \text{variables}$) : $\Rightarrow \text{values}$
bind($_ : \text{identifiers}, _ : \text{values}$) : $\Rightarrow \text{environments}$
bound($_ : \text{identifiers}$) : $\Rightarrow \text{values}$
checked($_ : (T)^?$) : $\Rightarrow T$
closed($_ : \Rightarrow T$) : $\Rightarrow T$
closure($_ : S \Rightarrow T$) : $\Rightarrow \text{abstractions}(S \Rightarrow T)$
collateral($_ : \text{environments}^*$) : $\Rightarrow \text{environments}$
control($_ : \text{functions}(\text{continuations}(T, S), S)$) : $\Rightarrow T$
effect($_ : T^*$) : $\Rightarrow \text{null-type}$
delimit-cc($_ : \Rightarrow T$) : $\Rightarrow T$
fail : $\Rightarrow \text{empty-type}$
give($_ : S, _ : S \Rightarrow T$) : $\Rightarrow T$
given : $T \Rightarrow T$
handle-return($_ : \Rightarrow T$) : $\Rightarrow T$
handle-thrown($_ : S \Rightarrow T, _ : X \Rightarrow T$) : $S \Rightarrow T$
hole : $\Rightarrow \text{values}$
if-true-else($_ : \text{booleans}, _ : \Rightarrow T, _ : \Rightarrow T$) : $\Rightarrow T$
left-to-right-repeat($_ : \text{integers} \Rightarrow T, _ : \text{integers}, _ : \text{integers}$) : $\Rightarrow (T)^*$
match($_ : \text{values}, _ : \text{values}$) : $\Rightarrow \text{environments}$
print($_ : \text{values}^*$) : $\Rightarrow \text{null-type}$
read : $\Rightarrow \text{values}$
resume-continuation($_ : \text{continuations}(S, T), _ : S$) : $\Rightarrow T$
return($_ : T$) : $\Rightarrow \text{empty-type}$
scope($_ : \text{environments}, _ : \Rightarrow T$) : $\Rightarrow T$
sequential($_ : (\Rightarrow \text{null-type})^*, \Rightarrow T$) : $\Rightarrow T$
stuck : $\Rightarrow \text{empty-type}$
throw($_ : T$) : $\Rightarrow \text{empty-type}$
while($_ : \Rightarrow \text{booleans}, _ : \Rightarrow \text{null-type}$) : $\Rightarrow \text{null-type}$

Appendix B.2. Value Operations

abstraction($_ : S \Rightarrow T$) : $\text{abstractions}(S \Rightarrow T)$
continuation($_ : \text{abstractions}(() \Rightarrow T)$) : $\text{continuations}(S, T)$
function($_ : \text{abstractions}(S \Rightarrow T)$) : $\text{functions}(S, T)$

index(_: nats, _ : values^{*}) : ⇒values[?]
integer-add(_: integers^{*}) : ⇒integers
integer-divide(_: integers, _ : integers) : ⇒integers[?]
integer-modulo(_: integers, _ : integers) : ⇒integers[?]
integer-multiply(_: integers^{*}) : ⇒integers
integer-negate(_: integers) : ⇒integers
integer-subtract(_: integers, _ : integers) : ⇒integers
is-in-set(_: T, _ : sets(T)) : ⇒booleans
is-in-type(_: values, _ : types) : ⇒booleans
is-less(_: integers, _ : integers) : ⇒booleans
is-less-or-equal(_: integers, _ : integers) : ⇒booleans
is-greater(_: integers, _ : integers) : ⇒booleans
is-greater-or-equal(_: integers, _ : integers) : ⇒booleans
is-equal(_: values, _ : values) : ⇒booleans
length(_: values^{*}) : ⇒nats
lookup(_: maps(S, T[?]), _ : S) : ⇒(T[?])[?]
map-override(_: maps(S, T)^{*}) : ⇒maps(S, T)
not(_: booleans) : ⇒booleans
null : null-type
pattern(_: abstractions(values⇒environments)) : patterns
returned(_: values) : returning
thrown(_: values) : throwing
tuple(_: T^{*}) : tuples(T^{*})
variable(_: locations, _ : types) : variables
vector(_: (T)^{*}) : vectors(T)
vector-elements(_: vectors(T)) : ⇒(T)^{*}