

FunKons: Component-Based Semantics in K

Peter D. Mosses and Ferdinand Vesely

Swansea University, Swansea, SA2 8PP, United Kingdom
p.d.mosses@swansea.ac.uk, csfvesely@swansea.ac.uk

Abstract. Modularity has been recognised as a problematic issue of programming language semantics, and various semantic frameworks have been designed with it in mind. Reusability is another desirable feature which, although not the same as modularity, can be enabled by it. The K Framework, based on Rewriting Logic, has good modularity support, but reuse of specifications is not as well developed.

The P_LanCompS project is developing a framework providing an open-ended collection of reusable components for semantic specification. Each component specifies a single fundamental programming construct, or ‘funcon’. The semantics of concrete programming language constructs is given by translating them to combinations of funcons. In this paper, we show how this component-based approach can be seamlessly integrated with the K Framework. We give a component-based definition of CinK (a small subset of C++), using K to define its translation to funcons as well as the (dynamic) semantics of the funcons themselves.

1 Introduction

Even very different programming languages often share similar constructs. Consider OCaml’s conditional ‘**if** E_1 **then** E_2 **else** E_3 ’ and the conditional operator ‘ $E_1 ? E_2 : E_3$ ’ in C. These constructs have different concrete syntax but similar semantics, with some variation in details. We would like to exploit this similarity when defining formal semantics for both languages by reusing parts of the OCaml specification in the C specification. With traditional approaches to semantics, reuse through “copy-paste-and-edit” is usually the only option that is available to us. By default, this is also the case with the K Framework [9,13]. This style of specification reuse is not systematic, and prone to error.

The semantic framework currently being developed by the P_LanCompS project¹ provides *fundamental constructs* (funcons) that address the issues of reusability in a systematic manner. Funcons are small semantic entities which express essential concepts of programming languages. These formally specified components can be composed to capture the semantics of concrete programming language constructs. A specification of Caml Light has been developed as an initial case study [3] and a case study on C# is in progress.

For example, the funcon `if-true` can be used to specify OCaml’s conditional expression. Semantics is given by defining a translation from the concrete construct to the corresponding funcon term:

¹ <http://www.plancomps.org/>

$$\llbracket \text{if } E_1 \text{ then } E_2 \text{ else } E_3 \rrbracket = \text{if-true}(\llbracket E_1 \rrbracket, \llbracket E_2 \rrbracket, \llbracket E_3 \rrbracket)$$

Since the conditional operator in C uses integer valued expressions as the condition, its translation will reflect this:

$$\llbracket E_1 ? E_2 : E_3 \rrbracket = \text{if-true}(\text{not}(\text{equal}(\llbracket E_1 \rrbracket, 0)), \llbracket E_2 \rrbracket, \llbracket E_3 \rrbracket)$$

We could also define an `if-non-zero` funcon that would match the C-conditional semantics exactly. However, the translation using `if-true` is so simple that there wouldn't be much advantage in doing so. We can reuse the `if-true` funcon, and with it, its semantic definition. This way, we also make the difference between the OCaml and C conditional construct transparent. Section 2 provides more information on funcons.

The project uses MSOS [10], a modular variant of structural operational semantics [11], to formally define individual funcons. However, the funcon approach can be seamlessly integrated with other sufficiently modular specification frameworks. In this paper we demonstrate how users of the K Framework might benefit from using funcons. We give a specification of CinK [8,9], a pedagogical subset of C++, by defining the translation using K's rewrite rules (Section 3). We also define the semantics of funcons involved in the translation and present some of them in Section 4. Section 5 offers an overview of related work and alternative approaches. We conclude and suggest directions of future work in Section 6.

The complete prototyped specification of CinK and K definitions of all funcons involved in the specification are available online.² Also included are CinK test programs which we have used to test our specification. Interested readers may run these programs themselves using the K tool. The original executable specification of CinK by Lucanu and Șerbănuță is available on GitHub.³

2 Fundamental Constructs

As mentioned in the Introduction, the PPlanCompS project is developing an open-ended collection of fundamental programming constructs, or 'funcons'. Many funcons correspond closely to simplified programming language constructs. However, each funcon has *fixed* syntax and semantics. For example, the funcon written `assign(E1, E2)` has the effect of evaluating `E1` to a variable, `E2` to a value (in any order), then assigning the value to the variable; it is well-typed only if `E1` is of type `variables(T)` and `E2` is of type `T`. In contrast, the language construct written '`E1 = E2`' may be interpreted as an assignment or as an equality test (and its well-typedness changes accordingly) depending on the language.

The syntax or *signature* of a funcon determines its name, how many arguments it takes (if any), the sort of each argument, and the sort of the result. The following *computation sorts* reflect fundamental conceptual and semantic distinctions in programming languages.

² <http://cs.swan.ac.uk/~csfvesely/wrla2014/>

³ <https://github.com/kframework/cink-semantic>

- The sort **Comm** (commands) is for funcons (such as `assign($E1$, $E2$)`) that are executed only for their effects; on normal termination, a command computes the fixed value `skip`.
- The sort **Expr** (expressions) is for funcons (such as `stored-value(E)` and `bound-value(I)`) that compute values of sort **Values**.
- The sort **Decl** (declarations) is for funcons (such as `bind-value(I , E)`) that compute environments of sort **Environments**, which represent sets of bindings between identifiers and values.

All computation sorts include their sorts of computed values as subsorts: a value takes no steps at all to compute itself.

One of the aims of the PLanCompS project is to establish an online repository of funcons (and data types) for anybody to use ‘off-the-shelf’ as components of language specifications. The project is currently testing the reusability of existing funcons and developing new ones in connection with some major case studies (including Caml Light, C#, and Java). Because individual funcons are meant to represent *fundamental* concepts in programming languages, many funcons (expressing, e.g., sequencing, conditionals, variable lookup and dereferencing) have a high potential for reuse. In fact, many funcons used in the Caml Light case study appear in the semantics of CinK presented in the following section.

The nomenclature and notation for the existing funcons are still evolving, and they will be finalised only when the case studies have been completed, in connection with the publication of the repository. Observant readers are likely to notice some (minor) differences between the funcon names used in this paper and in previous papers (e.g. [3]).

Regardless of the details of funcon notation, funcons can be algebraically composed to form funcon terms, according to their argument sorts (strictly lifted to corresponding computation sorts) and result sorts. *Well-formedness* of funcon terms is context-free: `assign($E1$, $E2$)` is a well-formed funcon term whenever $E1$ and $E2$ are well-formed funcon terms of sort **Expr**. In contrast, *well-typedness* of funcon terms is generally context-sensitive. For example, the funcon term `assign(bound-value(I), 42)` is well-typed only in the scope of a declaration that binds I to an integer variable. Dynamic semantics is defined for all well-formed terms; execution of ill-typed terms usually fails.

The composability of funcons does *not* depend on features such as whether they might have side effects, terminate abruptly, diverge, spawn processes, interact, etc. This is crucial for the reusability of the funcons. The semantics of each funcon has to be specified without regard to the context in which it might be used, which requires a highly modular specification framework. Funcon specifications have previously been given in MSOS, Rewriting Logic, ASF+SDF, and action notation. Here, we explore specifying funcons in K, following Roşu.⁴

A component-based semantics of a programming language is specified by a context-free grammar for an abstract syntax for the language, together with a family of inductively specified functions translating abstract syntax trees to

⁴ [k/examples/funcons](http://www.kframework.org) in the stable K distribution at <http://www.kframework.org>

funcon terms. The static and dynamic semantics of a program is given by that of the resulting funcon term. As mentioned above, funcons have fixed syntax and semantics. Thus, evolution of a language is expressed as changes to translation functions. If the syntax or semantics of the programming language changes, the definition of the translation function is updated to reflect this.

Tool support for translating programs to funcon terms, and for executing the static and dynamic semantics of such terms, has previously been developed in Prolog [2], Maude [1] and ASF+SDF. We now present our experiment with K, focusing on dynamic semantics.

3 CinF: a Funcon Specification of CinK

This section presents an overview of our CinK specification using funcons. We include examples from the K sources of the specification. A selection of definitions of funcons involved in the specification can be found in Section 4.

CinK is a pedagogical subset of C++ [8,9] used for experimentation with the K Framework. The original report [8] presents the language in seven iterations. The first specifies a basic imperative language; subsequent iterations extend it with threads, model-checking, references, pointers, and unidimensional and multi-dimensional arrays. Our specification starts with only an expression language which we extend with declarations, statements, functions, threads, references, and pointers. The extensions follow the order of the CinK iterations, however we omit arrays and support for model-checking.

The grammar which we have used for our specification is a simplified grammar matching CinK derived from the C++ grammar found in the standard [7, Appendix A].

We invite the reader to compare our specification by translation to funcons with the original K specification of CinK in [8]. Our hope is that our translation functions, together with the suggestive naming of funcons give a rough understanding of the semantics of language constructs, even before looking at the semantics of funcons themselves.

3.1 Simple Expressions

To give semantics for expressions we use the translation function `evaluate[[_]]` : `Expression` \rightarrow `Expr`. It produces a funcon term (of sort `Expr`) which, when executed, evaluates the argument expression.

Definitions for arithmetic expressions in CinK can be given very straightforwardly using data operations, which all extend to strict funcons on `Expr`. Numeric types in CinK are limited to integers with some common operations. For example, semantics of the multiplication operator is expressed as the application of the operation `int-times` to translations of operand expressions:

```
rule evaluate[[ E1:Expression * E2:Expression ]]  $\Rightarrow$ 
      int-times(evaluate[[ E1 ]], evaluate[[ E2 ]])
```

The ‘short-circuit and’ operator can be readily expressed using a conditional funcon, which is strict only in its first argument. The K definition for `if-true` can be found in Section 4.

```
rule evaluate[[ E1:Expression && E2:Expression ]] =>
  if-true(evaluate[[ E1 ]], evaluate[[ E2 ]], false)
```

We will use the generic `if-true` funcon later in this section to define the conditional statement.

3.2 Variables, Blocks and Scope

Bindings and Variables Semantics of declarations are given using `elaborate[[_]] : DeclarationSeq → Decl`. The `bind-value(I, V)` funcon binds the identifier I to the value V , producing a ‘small’ environment containing only the newly created binding. To allocate a new variable of a specified type we use `allocate`. In Caml Light, `bind-value` is used for individual name-value bindings in `let`-expressions and `allocate` for reference data types (e.g. “`ref int`”).

```
rule elaborate[[ T:TypeSpecifier I:Id ; ]] =>
  bind-value(I, allocate(variables(type[[ T ]])))
```

In relation to variables, CinK (following C++) distinguishes between two general categories of expressions: *lvalue*- and *rvalue*-expressions. We express this distinction by having different translation functions for expressions in `lvalue` and `rvalue` contexts: in addition to `evaluate[[_]]`, we define `evaluate-lval[[_]]` and `evaluate-rval[[_]]`. The default function `evaluate[[_]]` produces terms evaluating `lvalue` and `rvalue` expressions according to their category. When an expression is expected to evaluate to an *lvalue*, we use `evaluate-lval[[_]]`, which is undefined on `rvalue` expressions. When an *rvalue* is expected, we use `evaluate-rval[[_]]` which produces terms evaluating all expressions into `rvalues`. For `lvalue` expressions it returns the corresponding stored value, i.e., it serves as an `lvalue-to-rvalue` conversion.

The addition of variables also affects our translations of simple expressions and we need to update them. For example, numeric operations expect an `rvalue` and thus the operands are now translated using `evaluate-rval[[_]]`.

To get the variable bound to an identifier in the current environment we use `bound-value`. A variable is dereferenced using `stored-value`. The semantics for an identifier appearing in an `lvalue` or `rvalue` context is thus:

```
rule evaluate-lval[[ I:Id ]] => bound-value(I)
rule evaluate-rval[[ I:Id ]] => stored-value(evaluate-lval[[ I ]])
```

Blocks and Controlling Scope We distinguish between declaration statements and other statements within a block using funcons `scope` and `seq`. The funcon `scope(D, X)` evaluates X in the current environment overridden with the environment computed by D . A declaration statement within a block produces a new environment that is valid until the end of the block:

rule $\text{execute} \llbracket BD:\text{BlockDeclaration } SS:\text{StatementSeq} \rrbracket \Rightarrow$
 $\text{scope}(\text{elaborate} \llbracket BD \rrbracket, \text{execute} \llbracket SS \rrbracket)$

The function $\text{execute} \llbracket - \rrbracket : \text{StatementSeq} \rightarrow \text{Comm}$ translates statements to funcon commands.

For all other kinds of statements in a block we use the simple sequencing funcon $\text{seq}(C, X)$ which executes the command C for side effects and then executes X .

rule $\text{execute} \llbracket BS:\text{BlockStatement } SS:\text{StatementSeq} \rrbracket \Rightarrow$
 $\text{seq}(\text{execute} \llbracket BS \rrbracket, \text{execute} \llbracket SS \rrbracket)$

To accumulate multiple declarations into one environment we use the `accum` funcon. The funcon $\text{accum}(D1, D2)$ is similar to `scope`, except its result is the environment produced by elaborating declaration $D2$ and overriding the environment computed by $D1$ with it. This matches the semantics of a multi-variable declaration:

rule $\text{elaborate} \llbracket T:\text{TypeSpecifier } ID:\text{InitDeclarator},$
 $IDL:\text{InitDeclaratorList}; \rrbracket \Rightarrow$
 $\text{accum}(\text{elaborate} \llbracket T ID; \rrbracket, \text{elaborate} \llbracket T IDL; \rrbracket)$

Note that `accum` is strict in its first argument and so the correct order of evaluation is enforced.

Although Caml Light and CinK are quite different languages, all the funcons we needed so far are reused from [3].

3.3 Assignment and Control Statements

The basic construct for updating variables in CinK/C++ is the assignment expression ‘ $E_1 = E_2$ ’, where the expression E_1 is expected to evaluate to an lvalue, to which the rvalue of E_2 will be assigned. The value of the whole expression is the lvalue of E_1 . Semantics of assignment is a rather simple translation using the `assign-giving-variable` funcon:

rule $\text{evaluate-lval} \llbracket E1:\text{Expression} = E2:\text{Expression} \rrbracket \Rightarrow$
 $\text{assign-giving-variable}(\text{evaluate-lval} \llbracket E1 \rrbracket, \text{evaluate-rval} \llbracket E2 \rrbracket)$

The funcon `assign-giving-variable` is strict in both arguments but not sequentially, so the arguments are evaluated in an unspecified order. The funcon assigns the value given as its second argument to the variable given as its first argument and returns this variable as result.

CinK has boolean-valued conditions and the translations of while- and if-statements are trivial:

rule $\text{execute} \llbracket \text{while} (E:\text{Expression}) S:\text{Statement} \rrbracket \Rightarrow$
 $\text{while-true}(\text{evaluate-rval} \llbracket E \rrbracket, \text{execute} \llbracket S \rrbracket)$
rule $\text{execute} \llbracket \text{if} (E:\text{Expression}) S:\text{Statement} \rrbracket \Rightarrow$
 $\text{execute} \llbracket \text{if} (E) S \text{ else } \{ \} \rrbracket$
rule $\text{execute} \llbracket \text{if} (E:\text{Expression}) S1:\text{Statement} \text{ else } S2:\text{Statement} \rrbracket \Rightarrow$
 $\text{if-true}(\text{evaluate-rval} \llbracket E \rrbracket, \text{execute} \llbracket S1 \rrbracket, \text{execute} \llbracket S2 \rrbracket)$

3.4 Function Definition and Calling

We represent functions as *abstraction values* which wrap any computation as a value. An abstraction can be passed as a parameter, bound to an identifier, or stored like any other value. To turn a function term into an abstraction, we use the `abstraction` value constructor. The function `apply` applies an abstraction to a value and the abstraction may refer to the passed value using `given`. Multiple parameters can be passed as a tuple constructed via tuple value constructors.

A function call expression simply applies the abstraction to translated arguments:

```
rule evaluate-rval [[ E1:Expression ( E2:Expression ) ]] =>
  apply(evaluate-rval [[ E1 ]], evaluate-params [[ tuple(E2) ]])
```

At this stage the language only supports call-by-value semantics and so each parameter is evaluated to an *rvalue* before being passed to a function. The translation function `evaluate-params` (defined in terms of `evaluate-rval`) recurses through the parameter expressions and constructs a tuple.

```
rule evaluate-params [[ tuple(E1:Expression , E2:Expression) ]] =>
  tuple-prefix(evaluate-rval [[ E1 ]], evaluate-params [[ tuple(E2) ]])
rule evaluate-params [[ tuple(E:Expression) ]] =>
  tuple-prefix(evaluate-rval [[ E ]], tuple(.))
```

We have introduced the auxiliary abstract syntax `tuple(E)` to ensure that parameters separated by commas are not interpreted as a comma-operator expression.

We use *patterns* as translations of function parameters. Patterns themselves are abstractions which compute an environment when applied to a matching value. The pattern for passing a single parameter by value allocates a variable of the corresponding type and binds it to an identifier; then it assigns the parameter value to the variable and returns the resulting environment.

```
rule pattern [[ T:TypeSpecifier I:Id ]] =>
  abstraction(
    accum(bind-value(I, allocate(variables(type [[ T ]]))),
      decl-effect(assign(bound-value(I), given))))
```

Here we use the function `decl-effect(C)`, which allows using a command `C` as a declaration. It is an abbreviation for `seq(C, bindings(.))`.

Roughly, the semantics of a function definition is to allocate storage for an abstraction of the corresponding type, bind it to the function name, and use it to store an abstraction of the function body. Looking closer, the definition has to deal with some more details:

```
rule elaborate [[ T:TypeSpecifier I:Id ( PDL:ParameterDeclarationList )
  CS:CompoundStatement ]] =>
  decl-effect(assign(bound-value(I),
    close(abstraction(
      scope(match-compound(pattern-tuple [[ PDL ]], given),
        catch(seq(execute [[ CS ]], throw(variant(returned, null))),
          abstraction(original(returned, given))))))))
```

Within the abstraction we use `match-compound` to match the passed value against the pattern tuple constructed from individual parameter patterns. The translation of the function body is evaluated in the environment produced by this matching (`scope`). Since a return statement abruptly terminates a function returning a value, we represent return statements as exceptions containing a value tagged with the atom ‘returned’ and wrap the function body in a handler. The `catch` funcon catches the exception and the handling abstraction retrieves the value tagged with ‘returned’, making it the return value of the whole function. In case there was no return statement in the body of the function, we throw a ‘returned’ with `null`. Using `close` we form a closure of the abstraction with respect to the definition-time environment. This imposes static scopes for bindings.

As mentioned above, an explicit return statement translates to throwing a value tagged with ‘returned’. A parameterless return throws a `null`.

```
rule execute[[ return E:Expression; ]] ⇒
  throw(variant(returned, evaluate-rval[[ E ]]))
rule execute[[ return ; ]] ⇒ throw(variant(returned, null))
```

As a simple way of allowing self- and mutually recursive function definitions, we pre-allocate function variables and bind all function names declared at the top-level in a global environment using `evaluate-forwards[[_]]`. Then we combine this environment with the elaboration of full function definitions and other declarations. The `main` function is called in the scope of the global environment.

```
rule translate[[ DS:DeclarationSeq ]] ⇒
  scope(accum(elaborate-forwards[[ DS ]], elaborate[[ DS ]]),
    effect(apply(evaluate-rval[[ main ]], tuple(.))))
```

Because function identifiers are already bound when the full function definition is elaborated, the full definition only assigns the abstraction to the pre-allocated variable.

3.5 Threads

The second iteration in the original CinK report adds very basic thread support to the language. Spawning a thread in CinK mimics the syntax of using the `std::thread` class from the C++ standard library. However, instead of referring to the standard library, semantics is given to the construct directly.

```
rule elaborate[[ std::thread I1:Id ( I2:Id , E:Expression ) ; ]] ⇒
  decl-effect(effect(spawn(close(abstraction(evaluate[[ I2 (E) ]]]))))
```

The funcon `spawn(A)` creates a new thread in which the abstraction *A* will be applied. In our case the abstraction contains a function call corresponding to the parameters given to the thread constructor.

3.6 References

A reference in C++ is an alias for a variable, i.e., it introduces a new name for an already existing variable.

rule elaborate[[T :TypeSpecifier & I :Id = E :Expression]] \Rightarrow
 bind-value(I , evaluate-lval[[E]])

The expression E is expected to be an lvalue and we bind the resulting variable to identifier I . We are assuming that the input program is statically correct and thus the variable will have the right type.

A reference parameter pattern simply binds I to the given variable.

rule pattern[[T :TypeSpecifier & I :Id]] \Rightarrow
 abstraction(bind-value(I , given))

Before introducing references, we evaluated function parameters to an rvalue. Now the function evaluate-param[[$-$]] has to be *redefined* in terms of evaluate[[$-$]] instead of evaluate-rval[[$-$]]. Dereferencing is handled conditionally inside the parameter pattern.

rule pattern[[T :TypeSpecifier I :Id]] \Rightarrow
 abstraction(
 accum(bind-value(I , allocate(variables(type[[T]]))),
 decl-effect(assign(bound-value(I), current-value(given))))))

The funcon current-value dereferences its parameter if it is a variable (lvalue), otherwise returns the parameter itself.

3.7 Pointers

The last CinK extension that we consider is the addition of pointers to the language. Pointers are variables that hold addresses to other objects in memory. A pointer declaration allocates a new object for holding locations (variables in our terminology). Our semantics of declarations uses types to allocate storage and a pointer declaration complicates matters. Here we present a simplified version only supporting single-level indirection. The complete version has to deal with the notoriously complicated syntax for pointer declarations in C++.

rule elaborate[[T :TypeSpecifier * I :Id ;]] \Rightarrow
 accum(bind-value(I , allocate(variables(type[[T *]]))))
rule type[[T :TypeSpecifier *]] \Rightarrow pointers(type[[T]])

For our full specification of pointers, we refer the reader to the online material.

Explicit dereferencing of a pointer variable amounts to retrieving the value stored in the pointer. This value is the location to which the pointer is pointing. This is expressed in our translation:

rule evaluate-lval[[* E :Expression]] \Rightarrow
 stored-value(evaluate-rval[[E]])

A Note on Reuse The complete funcon definition of CinK available online uses 26 funcons. Of these, 19 have been previously used in the specification of Caml Light and only 7 were introduced in the present work, 3 of which are just abbreviations for longer funcon terms. It is thus possible to conclude that the reuse of funcons between the Caml Light and CinK specifications is high, even if the languages are quite different.

3.8 Configuration

The configuration of the final iteration of our specification is as follows:

```

configuration
  <T>
    <threads>
      <thread multiplicity="*">
        <name> main:Threads </name>
        <k> translate[[ $PGM:TranslationUnit ]] </k>
        <xstack> .List </xstack>
        <context>
          <env> .Map </env>
          <given> no-value </given>
        </context>
      </thread>
    </threads>
    <store> .Map </store>
    <output stream="stdout"> .List </output>
    <input stream="stdin"> .List </input>
  </T>

```

It appears that this configuration could be generated from the K rules defining the funcons used in our specification of CinK. It is unclear to us whether inference of K configurations from arbitrary K rules is possible, and whether it would be consistent with the K configuration abstraction algorithm.

3.9 Sequencing of Side Effects

Following the C++ standard [7], CinK decouples side effects of some constructs to allow delaying memory writes to after an expression value has been returned. This gives compilers more freedom for performing optimisations and during code generation. The newest C++ standard uses a relation *sequenced before* to define how side effects are to be ordered with respect to each other and to value evaluation. The CinK specification uses auxiliary constructs for side effects and uses a bag to collect side effects. An auxiliary sequence point construct forces finalisation of side effects in the bag. We are currently experimenting with funcons to express decoupled side effects.

4 Funcons in K

We now illustrate our K specification of the syntax and semantics of the funcons and value types used in our component-based analysis of CinF. We specify each funcon and value type in a separate module, to facilitate selective reuse. Since modularity is a significant feature of our specifications, we show some of the specified imports. The complete specifications are available online, together with the K specification of the translation of CinF programs to funcons.

4.1 Expressions

Expressions compute values:

```

module EXPR imports VALUES
  syntax Expr ::= Values
  syntax KResult ::= Values

```

Our specifications of value types lift the usual value operations to expression funcons, each of which is strict in all its arguments:

```

module INTEGERS imports EXPR ...
  syntax Expr ::= "int-times" "(" Expr "," Expr ")" [strict]
  | ...
  syntax Values ::= Int
  rule int-times(I1:Int, I2:Int)  $\Rightarrow$  I1 *Int I2
  rule ...

```

In contrast, the conditional expression funcon `if-true(E1, E2, E3)` is strict only in *E1*, and its rules involve unevaluated expression arguments:

```

module IF-TRUE-EXPR imports EXPR ...
  syntax Expr ::= "if-true" "(" Expr "," Expr "," Expr ")" [strict(1)]
  rule if-true(true, E:Expr, _)  $\Rightarrow$  E
  rule if-true(false, _, E:Expr)  $\Rightarrow$  E

```

We specify a corresponding funcon for conditional commands separately, since it appears that K modules cannot have parametric sorts (although the rules above could be generalised to arbitrary *K* arguments).

4.2 Declarations

```

module DECL imports BINDINGS
  syntax Decl ::= Bindings
  syntax KResult ::= Bindings

```

Bindings are values corresponding to environments (mapping identifiers to values), and come equipped with some operations that can be used to compose declarations:

```

module BINDINGS imports DECL
  syntax Bindings ::= bindings(Map)
  syntax Decl ::= "bindings-union" "(" Decl "," Decl ")" [strict]
  rule bindings-union(bindings(M1:Map), bindings(M2:Map))  $\Rightarrow$ 
    bindings(M1 M2)

```

We could have included the funcon `bind-value(I, E)` as an operation in the above module, since it is strict in its only expression argument:

```

module BIND-VALUE imports ...
  syntax Decl ::= "bind-value" "(" Id "," Expr ")" [strict(2)]
  rule bind-value(I:Id, V:Values)  $\Rightarrow$  bindings(I |-> V)

```

In contrast, the following funcons involve inspecting or (temporarily) changing the current environment, which is assumed to be in an accompanying cell:

```

module BOUND-VALUE imports ...
  syntax Expr ::= "bound-value" "(" Id ")"
  rule <k> bound-value(I:Id) ⇒ V:Values ...</k>
      <env>... I |-> V ...</env>

module SCOPE-COMM imports ...
  syntax Comm ::= "scope" "(" Decl "," Comm ")" [strict(1)]
  rule <k> scope(bindings(Env:Map), C:Comm) ⇒
      reset-env(Env', C) ...</k>
      <env> Env':Map ⇒ Env'[Env] </env>

module ACCUM imports ...
  syntax Decl ::= "accum" "(" Decl "," Decl ")" [strict(1)]
  rule <k> accum(bindings(Env:Map), D:Decl) ⇒
      reset-env(Env', bindings-union(bindings(Env), D)) ...</k>
      <env> Env':Map ⇒ Env'[Env] </env>

```

The auxiliary operation `reset-env(M, K)` preserves the result of `K` when resetting the current environment to `M`:

```

module RESET-ENV
  syntax K ::= "reset-env" "(" Map "," K ")" [strict(2)]
  rule <k> reset-env(Env:Map, V':KResult) ⇒ V' ...</k>
      <env> _:Map ⇒ Env </env>

```

The `K` argument could be of sort `Expr`, `Decl` or `Comm`. Since we do not use `reset-env` directly in the translation of `CinF` to funcons, the fact that `reset-env(M, K)` is (semantically) of the same sort as `K` is irrelevant.

4.3 Commands

```

module COMM imports SKIP
  syntax Comm ::= Skip
  syntax KResult ::= Skip

```

In contrast to the usual style in `K` specifications, commands compute the unique value `skip:Skip` on normal termination, rather than dissolving. However, this difference does not affect the translation of programs to funcons.

```

module SEQ-DECL imports ...
  syntax Decl ::= "seq" "(" Comm "," Decl ")" [strict(1)]
  rule seq(skip, D:Decl) ⇒ D

```

As with `if-true`, the funcon `seq(C, X)` is essentially generic in `X`, but its syntax needs to be specified separately for each sort of `X`. In contrast, the sort of `effect(X)` is independent of the sort of `X`, and we can specify it generically:

```

module EFFECT imports COMM
  syntax Comm ::= "effect" "(" K ")" [strict]
  rule effect(_:KResult) ⇒ skip

```

The specification of `while-true` illustrates reuse *between* funcon specifications:

```

module WHILE-TRUE
  imports COMM
  imports EXPR
  imports IF-TRUE-COMM
  imports SEQ-COMM
  syntax Comm ::= "while-true" "(" Expr "," Comm ")"
  rule while-true(E:Expr, C:Comm) ⇒
    if-true(E, seq(C, while-true(E, C)), skip)

```

4.4 Variables

Variables are themselves treated as values:

```

module VARIABLES imports ...
  syntax Variables ::= "no-variable"
  syntax Values ::= Variables

```

The specifications of the funcons for allocating, assigning to, and inspecting the values stored in variables are much as usual, and we omit them here.

4.5 Functions

```

module FUNCTIONS imports ...
  syntax Functions ::= "abstraction" "(" Expr ")"
  syntax Values ::= Functions

```

The operation `abstraction(E)` constructs a value from an *unevaluated* expression *E*. It can then be closed to obtain static bindings for identifiers in *E* (the K specification of the funcon `close(E)` is unsurprising, and omitted here).

```

module APPLY imports ...
  syntax Expr ::= "apply" "(" Expr "," Expr ")" [strict]
  rule apply(abstraction(E:Expr), V:Values) ⇒ supply(V, E)

```

The funcon `supply(E1, E2)` makes the value of *E1* available as ‘given’ in the evaluation of *E2*:

```

module SUPPLY-EXPR imports ...
  syntax Expr ::= "supply" "(" Expr "," Expr ")" [strict(1)]
  rule <k> supply(V:Values, E:Expr) ⇒ reset-given(V', E) ...</k>
    <given> V' ⇒ V </given>

module GIVEN imports ...
  syntax Expr ::= "given"
  rule <k> given ⇒ V:Values ...</k> <given> V </given>

```

The specifications of the funcons `throw` and `catch` assume that all cells used to represent the current context of a computation are grouped under a unique context cell. This gives improved modularity: the specification remains the same when further contextual cells are required. In other respects, the specification follows the usual style in the K literature, using a stack of exception handlers:

```

module THROW imports ...
syntax Comm ::= "throw" "(" Expr ")" [strict]
rule <k> (throw(V':Values) ~> _) => (apply(F, V') ~> K) </k>
      <xstack> (F:Functions, K:K, B:Bag) => . ...</xstack>
      <context> _ => B </context>

module CATCH imports ...
syntax Expr ::= "catch" "(" Comm "," Expr ")" [strict(2)]
rule <k> (catch(C:Comm, F:Functions) => (C ~> popx)) ~> K </k>
      <xstack> . => (F, K, B) ...</xstack>
      <context> B:Bag </context>
syntax K ::= "popx"
rule <k> popx => . </k> <xstack> _:ListItem => . ...</xstack>

```

Funcons `throw` and `catch` have the most complicated definitions of all, yet they are still modest in size and complexity.

5 Related Work

The work in this paper was inspired by a basic specification of the IMP example language in funcons using K by Roşu. IMP contains arithmetic and boolean expressions, variables, if- and while-statements, and blocks. The translation to funcons is specified directly using K rewrite rules without defining sorted translation functions. The example can be found in the stable K distribution.⁵

CinK, the sublanguage of C++ that we use as a case study in this paper, is taken from a technical report by Lucanu and Şerbănuţă [8]. We have limited ourselves to the same subset of C++, except that we omit arrays.

SIMPLE [12] is another K example language which is fairly similar to CinK. The language is presented in two variants: an untyped and a typed one. The definition of typed SIMPLE uses a different syntax and only specifies static semantics. With the component-based approach, we specify a single translation of language constructs to funcons. The MSOS of the funcons defines separate relations for typing and evaluation; in K, it seems we would need to provide a separate static semantics module for each funcon, since the strictness annotations and the computation rules are different.

K specifications scale up to real-world languages, as illustrated by Ellison's semantics of C [4]. The PPlanCompS project is currently carrying out major case studies (C#, Java) to examine how the funcon-based approach scales up to large languages, and to test the reusability of the funcon specifications.

Specification of individual language constructs in separate K modules was proposed by Hills and Roşu [6] and further developed by Hills [5, Chapter 5]. They obtained reusable rules by inferring the transformations needed for the rules to match the overall K configuration. The reusability of their modules was limited by their dependence on language syntax, and by the fact that the semantics of individual language constructs is generally more complicated than that of individual funcons.

⁵ <http://www.kframework.org>

6 Conclusion

We have given a component-based specification of CinK, using K to define the translation of CinK to funcons as well as the (dynamic) semantics of the funcons themselves. This experiment confirms the feasibility of integrating component-based semantics with the K Framework.

The K specification of each funcon is an independent module. Funcons are significantly simpler than constructs of languages such as CinK, and it was pleasantly straightforward to specify their K rules. However, we would have preferred the K configurations for combination of funcons to be generated automatically.

Many of the funcons used here for CinK were introduced in the component-based specification of Caml Light [3], demonstrating their reusability. The names of the funcons are suggestive of their intended interpretation, so the translation specification alone should convey a first impression of the CinK semantics. Readers are invited to browse the complete K specifications of our funcons online, then compare our translation of CinK to funcons with its direct specification in K [8].

In continuation of this work, we are investigating funcons to specify deferred side-effects and sequence points. We are also aiming to define the static semantics of funcons in K, so our translation would induce a static semantics for CinK.

References

1. F. Chalub and C. Braga. Maude MSOS tool. In *WRLA 2006*, volume 176 of *ENTCS*, pages 133 – 146. Elsevier, 2007.
2. M. Churchill and P. D. Mosses. Modular bisimulation theory for computations and values. In *FoSSaCS 2013*, volume 7794 of *LNCS*, pages 97–112. Springer, 2013.
3. M. Churchill, P. D. Mosses, and P. Torrini. Reusable components of semantic specifications. In *MODULARITY'14*. ACM, 2014. To appear.
4. C. Ellison and G. Roşu. An executable formal semantics of C with applications. In *POPL'12*, pages 533–544. ACM, 2012.
5. M. Hills. *A Modular Rewriting Approach to Language Design, Evolution and Analysis*. PhD thesis, University of Illinois at Urbana-Champaign, 2009.
6. M. Hills and G. Roşu. Towards a module system for K. In *WADT'08*, volume 5486 of *LNCS*, pages 187–205. Springer, 2009.
7. ISO International Standard ISO/IEC 14882:2011(E) – Programming Language C++, 2011. <http://isocpp.org/std/the-standard>.
8. D. Lucanu and T. F. Şerbănuţă. CinK – an exercise on how to think in K. Tech. Rep. TR 12-03 (V2), Faculty of Comput. Sci., A. I. Cuza Univ., Dec. 2013.
9. D. Lucanu, T. F. Şerbănuţă, and G. Roşu. K Framework distilled. In *WRLA'12*, volume 7571 of *LNCS*, pages 31–53. Springer, 2012.
10. P. D. Mosses. Modular structural operational semantics. *J. Log. Algebr. Program.*, 60-61:195–228, 2004.
11. G. D. Plotkin. A structural approach to operational semantics. *J. Log. Algebr. Program.*, 60-61:17–139, 2004.
12. G. Roşu and T. F. Şerbănuţă. K overview and SIMPLE case study. In *K'11*, ENTCS. Elsevier, 2014. To appear.
13. T. F. Şerbănuţă, A. Arusoaie, D. Lazar, C. Ellison, D. Lucanu, and G. Roşu. The K primer (v3.3). In *K'11*, ENTCS. Elsevier, 2014. To appear.