

# A DSL for Cross-Domain Security

David S. Hardin  
dshardin@rockwellcollins.com

Konrad L. Slind  
klsind@rockwellcollins.com

Rockwell Collins  
Advanced Technology Center

Michael W. Whalen  
mike.whelen@gmail.com

Tuan-Hung Pham  
hungpt43@gmail.com

University of Minnesota  
Software Engineering Center

## ABSTRACT

Guardol is a domain-specific language focused on the creation of high-assurance network guards and the specification of guard properties. The Guardol system generates Ada code from Guardol programs and also provides specification and automated verification support. Guard programs and specifications are translated to higher order logic, then deductively transformed to a form suitable for a SMT-style decision procedure for recursive functions over tree-structured data. The result is that difficult properties of Guardol programs can be proved fully automatically.

## Categories and Subject Descriptors

F.3.1 [Specification and Verifying and Reasoning about Programs]

## Keywords

Logics of programs, Mechanical verification, Specification techniques

## 1. INTRODUCTION

A *guard* is a device that mediates information sharing over a network between security domains according to a specified policy. Typical guard operations include reading field values in a packet, changing fields in a packet, transforming a packet by adding new fields, dropping fields from a packet, constructing audit messages, and removing a packet from a stream.

Guards are becoming prevalent, for example, in coalition forces networks, where selective sharing of data among coalition partners in real time is essential. One such guard, the Rockwell Collins Turnstile high-assurance, cross-domain guard [12], provides directional, bi-directional, and all-way guarding for up to three Ethernet connected networks. See Figure 1 for typical guard usage in a network. The proliferation of guards in critical applications, each with its own specialized language for specifying guarding functions, has

led to the need for a portable, high-assurance guard language.



Figure 1: Typical guard configuration

*Guardol* is a new, domain-specific programming language aimed at improving the creation, verification, and deployment of network guards. Guardol supports a wide variety of guard platforms, and features the ability to glue together existing or mandated functionality; the generation of both implementations and formal analysis artifacts; and sound, highly automated formal analysis.

Messages to be guarded, such as XML, may have recursive structure; thus a major aspect of Guardol is datatype declaration facilities similar to those available in functional languages such as SML [19] or Haskell [21]. Recursive programs over such datatypes are supported by ML-style pattern-matching. However, Guardol is not simply an adaptation of a functional language to guards. In fact, much of the syntax and semantics of Guardol is similar to that of Ada: Guardol is a sequential imperative language with non-side-effecting expressions, assignment, sequencing, conditional commands, and procedures with *in/out* variables. To a first approximation **Guardol** = **Ada** + **ML**. This hybrid language supports writing complex programs over complex data structures, while also providing standard programming constructs from Ada.

The Guardol system integrates several distinct components, as illustrated in Figure 2. A Guardol program in file `x.gdl` is parsed and typechecked by the Gryphon verification framework [18] developed by Rockwell Collins. Gryphon provides a collection of passes over Guardol ASTs that help simplify the program. From Gryphon, guard implementations can be generated—at present only in Ada—from Guardol descriptions. For the most part, this is conceptually simple since much of Guardol is a subset of Ada. However,




Figure 2: Guardol system components

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HILT'12, December 2–6, 2012, Boston, Massachusetts, USA.  
Copyright 2012 ACM 978-1-4503-1505-0/12/12 ...\$15.00.

datatypes and pattern matching need special treatment: the former requires automatic memory management, which we have implemented via a reference-counting style garbage collection scheme; while the latter requires a phase of pattern-match compilation [23]. Since our intent in this paper is mainly to discuss the verification path, we will omit further details.

## 2. AN EXAMPLE GUARD

In the following we will examine a simple guard written in Guardol. The guard applies a platform-supplied *dirty-word* operation DWO over a binary tree of messages (here identified with strings). When applied to a message, DWO can leave it unchanged, change it, or reject it with an audit string via `MsgAudit`.

```
type Msg      = string;
type MsgResult = {MsgOK : Msg | MsgAudit : string};
imported function DWO(Text : in Msg, Output : out MsgResult);
```

A `MsgTree` is a binary tree of messages. A `MsgTree` element can be a `Leaf` or a `Node`; the latter option is represented by a record with three fields. When the guard processes a `MsgTree` it either returns a new, possibly modified, tree, or returns an audit message.

```
type MsgTree = {Leaf
  | Node : [Value : Msg;
            Left : MsgTree; Right : MsgTree]};
type TreeResult = {TreeOK : MsgTree | TreeAudit : string};
```

The guard procedure takes its input tree in variable `Input`; and the return value, which has type `TreeResult`, is placed in `Output`. The body uses local variables for holding the results of recursing into the left and right subtrees, as well as for holding the result of calling `DWO`. The guard code is written as follows:

```
function Guard (Input : in MsgTree, Output : out TreeResult) =
begin
  var ValueResult : MsgResult;
      LeftResult, RightResult : TreeResult;
  in
  match Input with
  MsgTree'Leaf ⇒ Output := TreeResult'TreeOK(MsgTree'Leaf);
  MsgTree'Node node ⇒ begin
    DWO(node.Value, ValueResult);
    match ValueResult with
    MsgResult'MsgAudit A ⇒ Output := TreeResult'TreeAudit(A);
    MsgResult'MsgOK ValueMsg ⇒ begin
      Guard(node.Left, LeftResult);
      match LeftResult with
      TreeResult'TreeAudit A ⇒ Output := LeftResult;
      TreeResult'TreeOK LeftTree ⇒ begin
        Guard(node.Right, RightResult);
        match RightResult with
        TreeResult'TreeAudit A ⇒ Output := RightResult;
        TreeResult'TreeOK RightTree ⇒
          Output := TreeResult'TreeOK(MsgTree'Node
            [Value : ValueMsg,
              Left : LeftTree, Right : RightTree])
      end
    end
  end
end
```

The guard processes a tree by a pattern-matching style case analysis on the `Input` variable. There are several cases to consider. If `Input` is a leaf node, processing succeeds. This is accomplished by tagging the leaf with `TreeOK` and assigning to `Output`. Otherwise, if `Input` is an internal node

(`MsgTree'Node node`), the guard applies DWO to the message held at the node and recurses through the subtrees (recursive calls are marked with boxes). Complications arise from the fact that an audit may arise from a subcomputation and must be immediately propagated. The code essentially lifts the error monad of the external operation to the error monad of the guard. Note that, throughout, constructors include the name of their datatype; this allows constructor names to be reused.

### 2.1 Specifying guard properties

Many verification systems allow programs to be annotated with assertions. Under such an approach, a program may become cluttered with assertions and assertions may involve logic constructs. Since we wanted to avoid clutter and shield programmers, as much as possible, from learning the syntax of a logic language, we decided to express specifications using Guardol programs. The key language construct facilitating verification is the *specification* declaration: it presents some code to be executed, sprinkled with assertions (which are just boolean program expressions).

Following is the specification for our example guard. The code runs the guard on the tree  $t$ , putting the result in  $r$ , which is either a `TreeOK` element or an audit (`TreeAudit`). If the former, then the returned tree is named  $u$  and the encoded property `Guard_Stable`, described below, must hold on it. On the other hand, if  $r$  is an audit, this means the property is vacuously true.

```
spec Guard_Correct = begin
  var t : MsgTree; r : TreeResult;
  in if ∀(M : Msg). DWO.Idempotent(M) then begin
    Guard(t, r);
    match r with
    TreeResult'TreeOK u ⇒ check Guard_Stable(u);
    TreeResult'TreeAudit A ⇒ skip;
  else skip;
```

The guard code is essentially parameterized by an arbitrary policy (DWO) on how messages are treated. The correctness property simply requires that the result obeys the policy. In other words, suppose the guard is run on tree  $t$ , returning tree  $u$ . If DWO is run on every message in  $u$ , we expect to get  $u$  back unchanged, since all dirty words should have been scrubbed out in the passage from  $t$  to  $u$ . This property is a kind of idempotence, coded up in the function `Guard_Stable`. The success of the correctness proof depends on the assump-

```
function Guard_Stable (MT : in MsgTree) returns Output : bool =
begin
  var R : MsgResult;
  in
  match MT with
  MsgTree'Leaf ⇒ Output := true;
  MsgTree'Node node ⇒ begin
    DWO(node.Value, R);
    match R with
    MsgResult'MsgOK M ⇒ Output := (node.Value = M);
    MsgResult'MsgAudit A ⇒ Output := false;
    Output := Output and Guard_Stable(node.Left)
      and Guard_Stable(node.Right);
  end
end
```

tion that the external dirty-word operation is idempotent on messages, which can be expressed by a computation.

`DWO.Idempotent` calls `DWO` twice, and checks that the result of the second call is the same as the result of the

```

function DWO_Idempotent(M : in Msg) returns Output : bool = begin
  var R1, R2 : MsgResult;
  in
  DWO(M, R1);
  match R1 with
  | MsgResult'MsgOK M2 => begin
    DWO(M2, R2);
    match R2 with
    | MsgResult'MsgOK M3 => Output := (M2 = M3);
    | MsgResult'MsgAudit A => Output := false;
  | MsgResult'MsgAudit A => Output := true;

```

first call, taking into account audits. If the first call returns an audit, then there is no second call, so the idempotence property is vacuously true. On the other hand, if the first call succeeds, but the second is an audit, that means that the first call somehow altered the message into one provoking an audit, so idempotence is definitely false.

### 3. SEMANTIC TRANSLATION

Verification of guards is performed using HOL4 [24] and an SMT solver, such as OpenSMT [4], CVC4 [1], or Z3 [5]. First, a program is mapped into a formal AST in HOL4, where the operational semantics of Guardol have been defined. One can reason directly in HOL4 about programs using the operational semantics; unfortunately, such an approach has limited applicability, requiring expertise in the use of a higher order logic theorem prover. Instead, we would like to make use of the high automation offered by SMT systems. However there is an obstacle: current SMT systems do not understand operational semantics.<sup>1</sup> We surmount the problem in two steps. First, *decompilation into logic* [20] is used to deductively map properties of a program in an operational semantics to analogous properties over a mathematical function equivalent to the original program. This places us in the realm of proving properties of recursive functions operating over recursive datatypes, an undecidable setting in general. The second step is to implement a decision procedure for functional programming [25]. This procedure necessarily has syntactic limitations, but it is able to handle a wide variety of interesting programs and their properties fully automatically.

#### 3.1 Translating programs to HOL

First, any datatypes in the program are defined in HOL (every Guardol type can be defined as a HOL type). Thus, in our example, the types `MsgTree`, `MsgResult`, and `TreeResult` are translated directly to HOL datatypes. Recognizers and selectors, *e.g.*, `isMsgTree_Leaf` and `destMsgTree_Node`, for these types are automatically defined and used to translate pattern-matching statements in programs to equivalent if-then-else representations. Programs are then translated into HOL *footprint*<sup>2</sup> function definitions. If a program is recursive, then a recursive function is defined. (Defining recursive functions in higher order logic requires a termination proof; for our example termination is automatically proved.) Note that the HOL footprint function for the guard in Figure 3 is second order due to the external operator `DWO`: all exter-

<sup>1</sup>The main reason for this state of affairs: there is not one operational semantics because each programming language has a different semantics. Another reason: the decision problem for such theories is undecidable.

<sup>2</sup>Terminology lifted from the separation logic literature.

nals are held in a record `ext` which is passed as a function argument.

```

Guard ext Input =
  if isMsgTree_Leaf Input then
    TreeResult_TreeOK MsgTree_Leaf
  else let
    ValueResult = ext.DWO((destMsgTree_Node Input).Value)
  in
  if isMsgResult_MsgAudit ValueResult then
    TreeResult_TreeAudit
      (destMsgResult_MsgAudit ValueResult)
  else let
    LeftResult = Guard ext ((destMsgTree_Node Input).Left)
  in
  if isTreeResult_TreeAudit LeftResult then
    LeftResult else let
    RightResult = Guard ext ((destMsgTree_Node Input).Right)
  in
  if isTreeResult_TreeAudit RightResult then
    RightResult
  else
    TreeResult_TreeOK(MsgTree_Node
      [Value := destMsgResult_MsgOK ValueResult;
        Left := destTreeResult_TreeOK LeftResult;
        Right := destTreeResult_TreeOK RightResult]

```

Figure 3: Footprint function of example guard

Similar translations are made for `Guard_Stable` and also `DWO_Idempotent`.

#### 3.2 Guardol operational semantics

The operational semantics of Guardol (see Figure 4) describes program evaluation by an inductively defined judgement saying how statements alter the program state. The formula `STEPS  $\Gamma$  code  $s_1$   $s_2$`  says “evaluation of statement `code` beginning in state  $s_1$  terminates and results in state  $s_2$ ”. (Thus we are giving a so-called *big-step* semantics.) Note that  $\Gamma$  is an environment binding procedure names to procedure bodies. The semantics follows an approach taken by Norbert Schirmer [22], wherein he constructed a *generic* semantics for a large class of sequential imperative programs, and then showed how to specialize the generic semantics to a particular programming language (a subset of C, for him). Similarly, Guardol is another instantiation of the generic semantics.

Evaluation is phrased in terms of a *mode* of evaluation, which describes a computation state. A computation state is either in `Normal` mode, or in one of a set of abnormal modes, including `Abrupt`, `Fault`, and `Stuck`. Usually computation is in `Normal` mode. However, if a `Throw` is evaluated, then computation proceeds in `Abrupt` mode. If a `Guard` command returns `false`, the computation transitions into a `Fault` mode. Finally, if the `Stuck` mode is entered, something is wrong, *e.g.*, a procedure is called but there is no binding for it in  $\Gamma$ .

The rules in Figure 4 are conventional, with the exception of `withState` which provides a dynamic program-generation facility, used to model blocks and procedure calls.

#### 3.3 Decompilation

The work of Myreen [20] shows how to decompile assembly programs to higher order logic functions; we do the same here for Guardol, a high-level language. A decompilation theorem

$\frac{[Skip]}{STEPS \Gamma \text{Skip (Normal } s) \text{ (Normal } s)}$	$\frac{[Basic]}{STEPS \Gamma \text{(Basic } f) \text{ (Normal } s) \text{ (Normal } (f \ s))}$	
$\frac{[Seq] \quad \frac{STEPS \Gamma \ c_1 \text{ (Normal } s_1) \ s_2 \quad STEPS \Gamma \ c_2 \ s_2 \ s_3}{STEPS \Gamma \text{ (Seq } c_1 \ c_2) \text{ (Normal } s_1) \ s_3}}$	$\frac{[withState] \quad \frac{STEPS \Gamma \text{ (f } s_1) \text{ Normal } s_1) \ s_2}{STEPS \Gamma \text{ (withState } f) \text{ (Normal } s_1) \ s_2}}$	
$\frac{[Cond-True] \quad \frac{P(s_1) \quad STEPS \Gamma \ c_1 \text{ (Normal } s_1) \ s_2}{STEPS \Gamma \text{ (Cond } P \ c_1 \ c_2) \text{ (Normal } s_1) \ s_2}}$	$\frac{[Cond-False] \quad \frac{\neg P(s_1) \quad STEPS \Gamma \ c_2 \text{ (Normal } s_1) \ s_2}{STEPS \Gamma \text{ (Cond } P \ c_1 \ c_2) \text{ (Normal } s_1) \ s_2}}$	
$\frac{[Call] \quad \frac{M.p \in \text{Dom}(\Gamma) \quad \Gamma(M.p) = c \quad STEPS \Gamma \ c \text{ (Normal } s_1) \ s_2}{STEPS \Gamma \text{ (Call } M.p) \text{ (Normal } s_1) \ s_2}}$		
$\frac{[Call-NotFound] \quad \frac{M.p \notin \text{Dom}(\Gamma)}{STEPS \Gamma \text{ (Call } M.p) \text{ (Normal } s) \text{ Stuck}}}$		
$\frac{[withState] \quad \frac{STEPS \Gamma \text{ (f } s_1) \text{ Normal } s_1) \ s_2}{STEPS \Gamma \text{ (withState } f) \text{ (Normal } s_1) \ s_2}}$		
$\frac{[Fault-Sink]}{STEPS \Gamma \ c \text{ (Fault } f) \text{ (Fault } f)}$	$\frac{[Stuck-Sink]}{STEPS \Gamma \ c \text{ Stuck Stuck}}$	$\frac{[Abrupt-Sink]}{STEPS \Gamma \ c \text{ (Abrupt } s) \text{ (Abrupt } s)}$
$\frac{[While-True] \quad \frac{P(s_1) \quad STEPS \Gamma \ c \text{ (Normal } s_1) \ s_2 \quad STEPS \Gamma \text{ (While } P \ c) \ s_2 \ s_3}{STEPS \Gamma \text{ (While } P \ c) \text{ (Normal } s_1) \ s_3}}$		
$\frac{[While-False] \quad \frac{\neg P(s)}{STEPS \Gamma \text{ (While } P \ c) \text{ (Normal } s) \text{ (Normal } s)}}$		

Figure 4: Evaluation rules

$$\begin{aligned} &\vdash \forall s_1 \ s_2. \forall x_1 \dots x_k. \\ &\quad s_1.proc.v_1 = x_1 \wedge \dots \wedge s_1.proc.v_k = x_k \wedge \\ &\quad STEPS \Gamma \boxed{code} \text{ (Normal } s_1) \text{ (Normal } s_2) \\ &\quad \Rightarrow \\ &\quad \text{let } (o_1, \dots, o_n) = \boxed{f(x_1, \dots, x_k)} \\ &\quad \text{in } s_2 = s_1 \text{ with } \{proc.w_1 := o_1, \dots, proc.w_n := o_n\} \end{aligned}$$

essentially states that evaluation of *code* implements footprint function *f*. The antecedent  $s_1.proc.v_1 = x_1 \wedge \dots \wedge s_1.proc.v_k = x_k$  equates  $x_1 \dots x_k$  to the values of program variables  $v_1 \dots v_k$  in state  $s_1$ . These values form the input for the function *f*, which delivers the output values which are used to update  $s_1$  to  $s_2$ .<sup>3</sup> Presently, the decompilation theorem only deals with code that starts evaluation in a **Normal** state and finishes in a **Normal** state.

### 3.3.1 The decompilation algorithm

Now we consider how to prove decompilation theorems for Guardol programs. It is important to emphasize that *decompilation is an algorithm*. It always succeeds, provided that all footprint functions coming from the Guardol program have been successfully proved to terminate.

<sup>3</sup>In our modelling, a program state is represented by a record containing all variables in the program. The notation  $s.proc.v$  denotes the value of program variable  $v$  in procedure  $proc$  in state  $s$ . The **with**-notation represents record update.

Before specifications can be translated to goals, the decompilation theorem

$$\vdash \forall s_1 \ s_2. \dots STEPS \Gamma \boxed{Call(qid)} \text{ (Normal } s_1) \text{ (Normal } s_2) \Rightarrow \dots$$

is formally proved for each procedure *qid* in the program, relating execution of the code for procedure *qid* with the footprint function for *qid*.

Decompilation proofs are automated by forward symbolic execution of *code*, using an environment of decompilation theorems to act as summaries for procedure calls. Table 1 presents rules used in the decompilation algorithm. For the most part, the rules are straightforward. We draw attention to the **Seq**, **withState**, and **Call** rules. The **Seq** (sequential composition) rule conjoins the results of simpler commands and introduces an existential formula ( $\exists t. \dots$ ). However, this is essentially universal since it occurs on the left of the top-level implication in the goal; thus it can be eliminated easily and occurrences of *t* can thenceforth be treated as Skolem constants. Both blocks and procedure calls in the Guardol program are encoded using **withState**. An application of **withState** stays in the current state, but replaces the current code by new code computed from the current state. Finally, there are two cases with the **Call** (procedure call) rule:

- The call is not recursive. In this case, the decompilation theorem for *qid* is fetched from the decompilation environment and instantiated, so we can derive

$$\begin{aligned} &\text{let } (o_1, \dots, o_n) = f(x_1, \dots, x_k) \\ &\text{in } s_2 = s_1 \text{ with } \{qid.w_1 := o_1, \dots, qid.w_n := o_n\} \end{aligned}$$

**Table 1: Rewrite rules in the decompilation algorithm**

Condition	Rewrite rule
$code = \text{Skip}$	$\vdash \text{STEPS } \Gamma \text{ Skip } s_1 s_2 = (s_1 = s_2)$
$code = \text{Basic}(f)$	$\vdash \text{STEPS } \Gamma \text{ Basic } (f) (\text{Normal } s_1) (\text{Normal } s_2) = (s_2 = f s_1)$
$code = \text{Seq}(c_1, c_2)$	$\vdash \text{STEPS } \Gamma (\text{Seq}(c_1, c_2)) (\text{Normal } s_1) (\text{Normal } s_2) =$ $\exists t. \text{STEPS } \Gamma c_1 (\text{Normal } s_1) (\text{Normal } t) \wedge \text{STEPS } \Gamma c_2 (\text{Normal } t) (\text{Normal } s_2)$
$code = \text{Cond}(P, c_1, c_2)$	$\vdash \text{STEPS } \Gamma (\text{Cond}(P, c_1, c_2)) (\text{Normal } s_1) (\text{Normal } s_2) =$ $\text{if } P s_1 \text{ then } \text{STEPS } \Gamma c_1 (\text{Normal } s_1) (\text{Normal } s_2)$ $\quad \text{else } \text{STEPS } \Gamma c_2 (\text{Normal } s_1) (\text{Normal } s_2)$
$code = \text{withState } f$	$\vdash \text{STEPS } \Gamma (\text{withState } f) (\text{Normal } s_1) s_2 = \text{STEPS } \Gamma (f s_1) (\text{Normal } s_1) s_2$
$code = \text{Call } qid$	depend on whether the function is recursive or not

where  $f$  is the footprint function for procedure  $qid$ . We can now propagate the value of the function to derive state  $s_2$ .

- The call is recursive. In this case, an inductive hypothesis in the goal—which is a decompilation theorem for a recursive call, by virtue of our having inducted at the outset of the proof—matches the call, and is instantiated. We can again prove the antecedent of the selected inductive hypothesis, and propagate the value of the resulting functional characterization, as in the non-recursive case.

The decompilation algorithm starts by either inducting, when the procedure for which the decompilation theorem is being proved is recursive, or not (otherwise). After applying the rewrite rules, at the end of each program path, we are left with an equality between states. The proof of this equality proceeds essentially by applying rewrite rules for normalizing states (recall that states are represented by records).

### 3.4 Translating specifications

A Guardol specification is intended to set up a computational context—a state—and then assert that a property holds in that state. In its simplest form, a specification looks like

```
spec name = begin
  var decls
  in
    code;
    check property;
end
```

where *property* is a boolean expression. A specification declaration is processed as follows. First, suppose that execution of *code* starts normally in  $s_1$  and ends normally in  $s_2$ , *i.e.*, assume  $\text{STEPS } \Gamma \text{ code } (\text{Normal } s_1) (\text{Normal } s_2)$ . We want to show that *property* holds in state  $s_2$ . This could be achieved by reasoning with the induction principle for STEPS, *i.e.*, by using the operational semantics; however, experience has shown that this approach is labor-intensive. We instead opt to formally leverage the decompilation theorem for *code*, which asserts that reasoning about the STEPS-behavior of *code* could just as well be accomplished by reasoning about function  $f$ . Thus, formally, we need to show

$$\begin{aligned} & (\text{let } (o_1, \dots, o_n) = f(x_1, \dots, x_k) \\ & \text{in } s_2 = s_1 \text{ with } \{ \text{name.w}_1 := o_1, \dots, \text{name.w}_n := o_n \} \\ & \Rightarrow \text{property } s_2 \end{aligned}$$

Now we have a situation where the proof is essentially about how facts about  $f$ , principally its recursion equations and induction theorem, imply the property. The original goal has been freed—by sound deductive steps—from the program state and operational semantics. The import of this, as alluded to earlier, is that a wide variety of proof tools become applicable. Interactive systems exemplified by ACL2, PVS, HOL4, and Isabelle/HOL have extensive lemma libraries and reasoning packages tailored for reasoning about recursively defined mathematical functions. SMT systems are also able to reason about such functions, via universal quantification, or by decision procedures, as we discuss in Section 4.

The simple form of specification above is not powerful enough to state many properties. Quite often, a collection of constraints needs to be placed on the input variables, or on external functions. To support this, specification statements allow checks sprinkled at arbitrary points in *code*:

```
spec name =
  begin locals
  in
    code[check P1, ..., check Pn]
  end
```

We support this with a program transformation, wherein occurrences of **check** are changed into assignments to a boolean variable. Let  $V$  be a boolean program variable not in *locals*. The above specification is transformed into

```
spec name =
  begin
    locals; V : bool;
  in
    V := true;
    code[V := V ∧ P1, ..., V := V ∧ Pn];
    check(V);
  end
```

Thus  $V$  is used to accumulate the results of the checks that occur throughout the code. Every property  $P_i$  is checked in the state holding just before the occurrence of  $\text{check}(P_i)$ , and all the checks must hold. This gives a flexible and concise way to express properties of programs, without embedding assertions in the source code of the program.

Recall the *Guard\_Correct* specification. Roughly, it says *If running the guard succeeds, then running Guard\_Stable on the result returns true*. Applying the decompiler to the code of the specification and using the resulting theorem to map from the operational semantics to the functional interpretation, we obtain the goal

$$\left( \begin{aligned} & (\forall m. \text{DWO.Idempotent } \text{ext } m) \wedge \\ & \text{Guard } \text{ext } t = \text{TreeResult\_TreeOK } t' \end{aligned} \right) \Rightarrow \text{Guard\_Stable } \text{ext } t'$$

which has the form required by our SMT prover, namely that the catamorphism `Guard_Stable` is applied to the result of calling `Guard`. However, an SMT prover may still not prove this goal, since the following steps need to be made: (1) inducting on the recursion structure of `Guard`, (2) expanding (once) the definition of `Guard`, (3) making higher order functions into first order, and (4) elimination of universal quantification.<sup>4</sup>

To address the first two problems, we induct with the induction theorem for `Guard`, which is automatically proved by HOL4, and expand the definition of `Guard` one step in the resulting inductive case. Thus we stop short of using the inductive hypotheses! The SMT solver will do that. The elimination of higher order functions is simple in the case of `Guard` since the function arguments (*ext* in this case) are essentially fixed constants whose behavior is constrained by hypotheses. This leaves the elimination of the universals; only the quantification on *m* in  $\forall m. \text{DWO\_Idempotent } \text{ext } m$  is problematic. We find all arguments of applications of `ext.DWO` in the body of `Guard`, and instantiate *m* to all of them (there’s only one in this case), adding all instantiations as hypotheses to the goal.

## 4. PROVING VERIFICATION CONDITIONS

The formulas generated as verification conditions from the previous section pose a fundamental research challenge: reasoning over the structure and contents of inductive datatypes. We have addressed this challenge through the use of a novel decision procedure recently developed by Suter, Dotta, and Kuncak [25]. This decision procedure (we call it SDK) can be integrated into an SMT solver to solve a variety of properties over recursive datatypes. It uses catamorphisms to create abstractions of the contents of tree-structured data that can then be solved using standard SMT techniques. The benefit of this decision procedure over other techniques involving quantifiers is that it is *complete* for a large class of reasoning problems involving datatypes, as described below.

### 4.1 Catamorphisms

In many reasoning problems involving recursive datatypes, we are interested in *abstracting* the contents of the datatype. To do this, we could define a function that maps the structure of the tree into a value. This kind of function is called a catamorphism [17] or *fold* function, which ‘folds up’ information about the data structure into a single value. The simplest abstraction that we can perform of a data structure is to map it into a Boolean result that describes whether it is ‘valid’ in some way. This approach is used in the function `Guard_Stable` in Section 2. We could of course create different functions to summarize the tree elements. For example, a tree can be abstracted as a number that represents the sum of all nodes, or as a tuple that describes the minimum and maximum elements within the tree. As long as the catamorphism is *sufficiently surjective* [25] and maps into a decidable theory, the procedure is theoretically *complete*. Moreover, we have found it to be *fast* in our initial experiments.

<sup>4</sup>Some of these steps are incorporated in various SMT systems, *e.g.*, many, but not all, SMT systems heuristically instantiate quantifiers. For a discussion of SMT-style induction see [15].

## 4.2 Overview of the Decision Procedure

The input of the decision procedure is a formula  $\phi$  of literals over elements of tree terms and tree abstractions ( $\mathcal{L}_C$ ) produced by the catamorphisms. The logic is *parametric* in the sense that we assume a datatype to be reasoned over, a catamorphism used to abstract the datatype, and the existence of a decidable theory  $C$  that is the result type of the catamorphism function. The syntax of the parametric logic is depicted in Figure 5.

The syntax of the logic ranges over datatype terms ( $T$  and  $S$ ), terms of a decidable collection theory  $C$ . Tree and collection theory formulas  $F_T$  and  $F_C$  describe equalities and inequalities over terms. The collection theory describes the result of catamorphism applications.  $E$  defines terms in the element types contained within the branches of the datatypes.  $\phi$  defines conjunctions of (restricted) formulas in the tree and collection theories. The  $\phi$  terms are the ones solved by the SDK procedure; these can be generalized to arbitrary propositional formulas ( $\psi$ ) through the use of a DPLL solver [8] which manages the other operators within the formula.

$S$	::=	$T \mid E$	Constructor args
$T$	::=	$t \mid \mathbb{C}_j(S_1, \dots, S_{n_j}) \mid \mathbb{S}_{j,k_\tau}(T)$	Tree terms
$C$	::=	$c \mid \alpha(T) \mid \mathcal{T}_C$	$C$ -terms
$F_T$	::=	$T = T \mid T \neq T$	Tree (dis)equalities
$F_C$	::=	$C = C \mid \mathcal{F}_C$	Formula of $\mathcal{L}_C$
$E$	::=	variables of type $\mathcal{E}_k \mid \mathbb{S}_{j,k_\mathcal{E}}(T)$	Expression
$\phi$	::=	$\bigwedge F_T \wedge \bigwedge F_C$	Conjunctions
$\psi$	::=	$\phi \mid \neg\phi \mid \phi \vee \phi \mid \phi \wedge \phi \mid \phi \Rightarrow \phi \mid \phi \Leftrightarrow \phi$	Formulas

Figure 5: Syntax of the parametric logic

In the procedure, we have a single datatype  $\tau$  with  $m$  constructors. The  $j$ -th constructor ( $1 \leq j \leq m$ ),  $\mathbb{C}_j$ , has  $n_j$  arguments ( $n_j \geq 0$ ), whose types are either  $\tau$  or  $\mathcal{E}$ , an element type. For each constructor  $\mathbb{C}_j$ , we have a list of selectors  $\mathbb{S}_{j,k}$  ( $1 \leq k \leq n_j$ ), which extracts the  $k$ -th argument of  $\mathbb{C}_j$ . For type safety, we may put the type of the argument to be extracted as a subscript of its selector. That is, each selector may be presented as either  $\mathbb{S}_{j,k_\tau}$  or  $\mathbb{S}_{j,k_\mathcal{E}}$ . The decision procedure is parameterized by  $\mathcal{E}$ , a collection type  $\mathcal{C}$ , and a catamorphism function  $\alpha : \tau \rightarrow \mathcal{C}$ . For example, the datatype `MsgTree` has two constructors `Leaf` and `Node`. The former has no argument while the latter has three arguments corresponding to its `Value`, `Left`, and `Right`. As a result, we have three selectors for `Node`, including `Value : MsgTree  $\rightarrow$  Msg`, `Left : MsgTree  $\rightarrow$  MsgTree`, and `Right : MsgTree  $\rightarrow$  MsgTree`. In addition, a tree can be abstracted by the catamorphism `Guard_Stable : MsgTree  $\rightarrow$  bool` to a boolean value. In the example,  $\mathcal{E}$ ,  $\mathcal{C}$ , and  $\alpha$  are `Msg`, `bool`, and `Guard_Stable`, respectively.

### 4.3 Implementing Suter-Dotta-Kuncak

The decision procedure works on top of an SMT solver  $\mathcal{S}^5$  that supports theories for  $\tau, \mathcal{E}, \mathcal{C}$ , and uninterpreted functions. Note that the only part of the parametric logic that is not inherently supported by  $\mathcal{S}$  is the applications of the catamorphism. Therefore, the main idea of the decision procedure is to approximate the behavior of the catamorphism

<sup>5</sup>Suter et al. [26] specifically used Z3 [5] as the underlying SMT solver.

by repeatedly unrolling it a certain number of times and treating the calls to the not-yet-unrolled catamorphism instances at the lowest levels as calls to uninterpreted functions. However, an uninterpreted function can return any values in its co-domain; hence, the presence of these uninterpreted functions can make the `sat/unsat` result not trustworthy. To address this issue, each time the catamorphism is unrolled, a boolean control condition  $B$  is created to determine if the uninterpreted functions at the bottom level are necessary to the determination of satisfiability. That is, if  $B$  is true, the list of uninterpreted functions does not play any role in the satisfiability result.

The main steps of the procedure are shown in Algorithm 1. The input of the algorithm is a formula  $\phi$  written in the parametric logic and a program  $\Pi$ , which contains  $\phi$  and the definitions of data type  $\tau$  and catamorphism  $\alpha$ . The goal of the algorithm is to determine the satisfiability of  $\phi$  through repeated unrolling  $\alpha$  using the `unrollStep` function. Given a formula  $\phi_i$  generated from the original  $\phi$  after unrolling the catamorphism  $i$  times and the corresponding control condition  $B_i$  of  $\phi_i$ , function `unrollStep`( $\phi_i, \Pi, B_i$ ) unrolls the catamorphism one more time and returns a pair  $(\phi_{i+1}, B_{i+1})$  containing the unrolled version  $\phi_{i+1}$  of  $\phi_i$  and a control condition  $B_{i+1}$  for  $\phi_{i+1}$ . Function `decide`( $\varphi$ ) simply calls  $S$  to check the satisfiability of  $\varphi$  and returns `SAT/UNSAT` accordingly.

```

1 ( $\phi, B$ )  $\leftarrow$  unrollStep( $\phi, \Pi, \emptyset$ )
2 while true do
3   switch decide( $\phi \wedge \bigwedge_{b \in B} b$ ) do
4     case SAT
5       | return “SAT”
6     case UNSAT
7       | switch decide( $\phi$ ) do
8         | case UNSAT
9           | return “UNSAT”
10        | case SAT
11          | ( $\phi, B$ )  $\leftarrow$  unrollStep( $\phi, \Pi, B$ )

```

**Algorithm 1:** Catamorphism unrolling algorithm [26]


Let us examine how satisfiability and unsatisfiability are determined in the algorithm. In general, the algorithm keeps unrolling the catamorphism until we find a `sat/unsat` result that we can trust. To do that, we need to consider several cases after each unrolling is carried out. First, at line 4,  $\phi$  is satisfiable and the control condition is true, which means uninterpreted functions are not involved in the satisfiable result. In this case, we have a complete tree model for the `SAT` result and we can conclude that the problem is satisfiable.

On the other hand, let us consider the case when `decide`( $\phi \wedge \bigwedge_{b \in B} b$ ) = `UNSAT`. The `UNSAT` may be due to the unsatisfiability of  $\phi$ , or the control condition, or both of them together. As a result, to understand the `UNSAT` more deeply, we could try to check the satisfiability of  $\phi$  alone, as depicted at line 7. Note that checking  $\phi$  alone also means that the control condition is not used; consequently, the values of uninterpreted functions may contribute to the `sat/unsat` result of `decide`( $\phi$ ). If `decide`( $\phi$ ) = `UNSAT` as at line 8, we can conclude that the problem is unsatisfiable because assigning

the uninterpreted functions to any values in their co-domains still cannot make the problem satisfiable as a whole. Finally, we need to consider the case `decide`( $\phi$ ) = `SAT` as at line 10. Since we already know that `decide`( $\phi \wedge \bigwedge_{b \in B} b$ ) = `UNSAT`, the only way to make `decide`( $\phi$ ) be `SAT` is by calling to at least one uninterpreted function, which also means that the `SAT` result is untrustworthy. Therefore, we need to keep unrolling at least one more time as denoted at line 11.

## 4.4 Tool Architecture

The overall architecture of our solver, called RADA (Reasoning over Abstract Datatypes with Abstraction), is shown in Figure 6. It closely follows the algorithm described in the previous section. We use CVC4 [1] and Z3 [5] as the underlying SMT solvers in RADA because of their powerful abilities to reason about recursive data types. The grammar of RADA in Figure 4.4 is based on the SMT-Lib 2.0 [2] format with some new syntax for selectors, testers, data type declarations, and catamorphism declarations.



**Figure 6:** RADA architecture.

$\langle \text{datadec} \rangle$	::=	( <b>declare-datatypes</b> ( $\langle \text{datatype} \rangle^+$ ) )
$\langle \text{datatype} \rangle$	::=	( ( $\langle \text{symbol} \rangle$ ) ( $\langle \text{datatype\_branch} \rangle^+$ ) )
$\langle \text{datatype\_branch} \rangle$	::=	( ( $\langle \text{symbol} \rangle$ ) [ ( $\langle \text{branch\_parameter} \rangle^+$ ) ] )
$\langle \text{branch\_parameter} \rangle$	::=	( ( $\langle \text{symbol} \rangle$ ) ( $\langle \text{sort} \rangle$ ) )
$\langle \text{catadec} \rangle$	::=	( <b>define-catamorphism</b> ( $\langle \text{cata} \rangle$ ) )
$\langle \text{cata} \rangle$	::=	( ( $\langle \text{symbol} \rangle$ ) ( ( $\langle \text{sort} \rangle$ ) ) ( $\langle \text{term} \rangle$ ) )
$\langle \text{selector\_application} \rangle$	::=	( $\langle \text{symbol} \rangle$ ) ( $\langle \text{symbol} \rangle$ )
$\langle \text{tester\_application} \rangle$	::=	<i>is_</i> ( $\langle \text{symbol} \rangle$ ) ( $\langle \text{symbol} \rangle$ )

**Figure 7:** RADA grammar.

Note that although selectors, testers, and data type declarations are not defined in SMT-Lib 2.0, all of them are currently supported by both CVC4 and Z3; therefore, only catamorphism declarations are not understood by these solvers. As a result, to bridge the gap between the input format of RADA and that of CVC4/Z3, each time the catamorphism is unrolled, we build an abstract syntax tree in which the catamorphism declaration is replaced by an uninterpreted function representing the behaviors of the unrolled parts of the catamorphism. Based on the abstract syntax tree, we generate an `.smt2` file that CVC4 or Z3 accepts with the help of a tool-specific emitter, which is responsible for creating a suitable `.smt2` file for the solver being used.

## 4.5 Experimental Results

To test our approach, we have developed a handful of small benchmark guard examples. For timings, we focus on the SMT solver which performs the interesting part of the proof search. The results are shown in Table 2, where the last guard is the running example from the paper. In our early experience, the SDK procedure allows a wide variety of interesting properties to be expressed and our initial timing experiments have been very encouraging, despite the relatively naïve implementation of SDK, which we will optimize

in the near future. For a point of comparison, we provide a translation of the problems to Microsoft’s Z3 in which we use universal quantifiers to describe the catamorphism behavior. This approach is incomplete, so properties that are falsifiable (i.e., return SAT) often do not terminate (we mark this as ‘unknown’). A better comparison would be to run our implementation against the Leon tool suite developed at EFPL [26]. Unfortunately, it is not currently possible to do so as the Leon tool suite operates directly over the Scala language rather than at the SMT level. All benchmarks were run on Windows 7 using an Intel Core I3 running at 2.13 GHz.

## 5. DISCUSSION

As a programming language with built-in verification support, Guardol seems amenable to being embedded in an IVL (Intermediate Verification Language) such as Boogie [16] or Why [7]. However, our basic aims would not be met in such a setting. The operational semantics in Section 3.2 defines Guardol program execution, which is the basis for verification. We want a strong formal connection between a program plus specifications, both expressed using that semantics, and the resulting SMT goals, which do not mention that semantics. The decompilation algorithm achieves this connection via machine-checked proof in higher order logic. This approach is simpler in some ways than an IVL, where there are two translations: one from source language to IVL and one from IVL to SMT goals. To our knowledge, IVL translations are not machine-checked, which is problematic for our applications. Our emphasis on formal models and deductive transformations should help Guardol programs pass the stringent certification requirements imposed on high-assurance guards.

Higher order logic plays a central role in our design. HOL4 implements a foundational semantic setting in which the following are formalized: program ASTs, operational semantics, footprint functions (along with their termination proofs and induction theorems), and decompilation theorems. Decompilation extensively uses higher order features when composing footprint functions corresponding to sub-programs. Moreover, the backend verification theories of the SMT system already exist in HOL4. This offers the possibility of doing SMT proof reconstruction [3] in order to obtain even higher levels of assurance. Another consequence is that, should a proof fail or take too long in RADA, it could be performed interactively.

An interesting issue concerns specifications: guard specifications can be about *intensional* aspects of a computation, e.g., its structure or sequencing, as well as its result. For example, one may want to check that data fuzzing operations always occur before encryption. However, our current framework, which translates programs to *extensional* functions, will not be able to use SMT reasoning on intensional properties. Information flow properties [9] are also intensional, since in that setting one is not only concerned with the value of a variable, but also whether particular inputs and code structures were used to produce it. Techniques similar to those in [27] could be used to annotate programs in order to allow SMT backends to reason about intensional guard properties.

### 5.1 Current and Future Work

We plan to further enhance both language and verification

aspects of Guardol. In recent work, we have built a custom Guardol editor, using the Xtext framework [6]. A screenshot is given in Figure 8. The editor provides conventional support such as keyword highlighting, code completion, and simple semantic processing, such as type-checking. The editor interface also supports the invocation of code generation and verification subsystems.

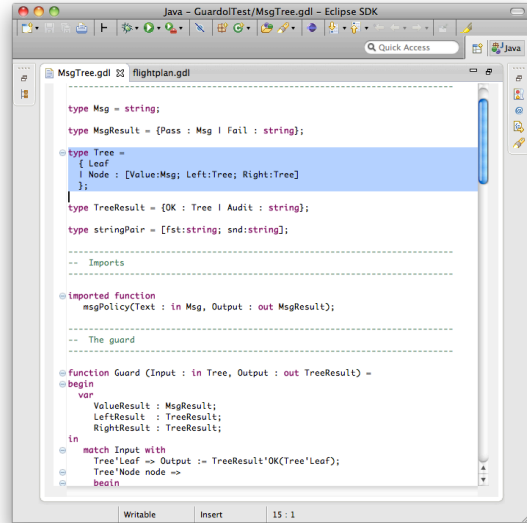


Figure 8: Xtext-based Guardol editor

The current Guardol language could be made more user-friendly. It is essentially a monomorphic version of ML with second order functions, owing to Guardol’s external function declarations. It might be worthwhile to support polymorphic types so that the repeated declarations of instances of polymorphic types, e.g., option types and list types, can be curtailed. Additionally, programs could be considerably more terse if exceptions were in the language, since explicitly threading the error monad wouldn’t be needed to describe guard failures.

Planned verification improvements include integration of string solvers and termination deferral. In the translation to SMT, strings are currently treated as an uninterpreted type and string operations as uninterpreted functions. Therefore, the system cannot reason in a complete way about guards where string manipulation is integral to correctness. We plan to integrate a string reasoner, e.g., [13], into our system to handle this problem. The use of string solvers in *web sanitizers* used to defeat cross-site scripting attacks [11] is highly similar to planned applications of Guardol. Moreover, guards and sanitizers have similar requirements, such as idempotence and commutativity.

A weak point in the end-to-end automation of Guardol verification is termination proofs. If the footprint function for a program happens to be recursive, its termination proof may well fail, thus stopping processing. Techniques for defining partial recursive functions have been developed [10, 14]. These approaches allow the derivation of recursion equations and induction theorems constrained by termination requirements. Adopting one of these techniques would remedy this weakness in Guardol, allowing the deferral of termination ar-



**Table 2: Experimental results on guard examples**

Test #	OpenSMT-SDK	Z3
Guard 1	5 sat (0.83s) / 1 unsat (0.1s)	5 unknown / 1 unsat (0.06 s)
Guard 2	3 unsat (0.28s)	1 unknown / 2 unsat (0.11 s)
Guard 3	3 sat (0.33s) / 4 unsat (0.46s)	1 unknown / 2 sat (0.17s) / 4 unsat (0.42 s)
DWO	1 unsat (0.34s)	1 unsat (0.11s)
MIME	17 unsat (2.0s)	17 unsat (0.2s)

guments while partial correctness proofs are addressed. Of course, termination arguments then need to be automated and we hope to apply RADA to useful subsets of this class of problems.

Arrays are an essential component in any programming language, and Guardol has arrays as a basic type. However, naive approaches to automating proofs of programs that iterate over arrays don't work well. We plan to identify a useful class of iterations over arrays adaptable to the SDK approach. The first obstacle is that general while-loops are not catamorphic. However, we intend to exploit the fact that the class of for loops is primitive recursive, and so every for-loop should have a catamorphic counterpart.

Finally, we plan to verify the correctness of code generated from guards. The general approach will be to relate the footprint function of a guard with the function resulting from decompiling the generated assembly code. Myreen's technology for decompilation of low-level code will be key in this effort.

## 6. REFERENCES

- [1] C. Barrett, C. L. Conway, M. Deters, L. Hadarean, D. Jovanović, T. King, A. Reynolds, and C. Tinelli. CVC4. In *Proceedings of the 23rd international conference on Computer aided verification, CAV'11*, pages 171–177, 2011.
- [2] C. Barrett, A. Stump, and C. Tinelli. The SMT-LIB standard: Version 2.0. In A. Gupta and D. Kroening, editors, *Proceedings of the 8th International Workshop on Satisfiability Modulo Theories (Edinburgh, England)*, 2010.
- [3] S. Böhme, A. Fox, T. Sewell, and T. Weber. Reconstruction of Z3's bit-vector proofs in HOL4 and Isabelle/HOL. In *Proceedings of Certified Programs and Proofs*, volume 7086 of *LNCS*. Springer, 2011.
- [4] R. Bruttomesso, E. Pek, N. Sharygina, and A. Tsitovich. The OpenSMT solver. In *Proceedings of TACAS*, volume 6015 of *LNCS*, 2010.
- [5] L. De Moura and N. Bjørner. Z3: An Efficient SMT Solver. In *Proceedings of the Theory and practice of software, 14th international conference on Tools and algorithms for the construction and analysis of systems, TACAS'08/ETAPS'08*, pages 337–340, 2008.
- [6] M. Eysholdt and H. Behrens. Xtext: implement your language faster than the quick and dirty way. In *Proceedings of the ACM international conference companion on Object oriented programming systems languages and applications companion, SPLASH '10*, pages 307–309. ACM, 2010.
- [7] J.-C. Filliâtre. *Deductive Program Verification*. Thèse d'habilitation, Université Paris 11, Dec. 2011.
- [8] H. Ganzinger, G. Hagen, R. Nieuwenhuis, A. Oliveras, and C. Tinelli. DPLL( $T$ ): Fast decision procedures. In *Proceedings of CAV*, volume 3114 of *LNCS*, pages 175–188. Springer, 2004.
- [9] J. Goguen and J. Meseguer. Security policies and security models. In *Proc of IEEE Symposium on Security and Privacy*, pages 11–20. IEEE Computer Society Press, 1982.
- [10] D. Greve. Assuming termination. In *Proceedings of ACL2 Workshop, ACL2 '09*, pages 114–122. ACM, 2009.
- [11] P. Hooimeijer, B. Livshits, D. Molnar, P. Saxena, and M. Veanes. Fast and precise sanitizer analysis with BEK. In *Proceedings of the 20th USENIX conference on Security*, pages 1–16, Berkeley, CA, USA, 2011. USENIX Association.
- [12] R. C. Inc. Turnstile High Assurance Guard Homepage. <http://www.rockwellcollins.com/>.
- [13] A. Kiezun, V. Ganesh, P. Guo, P. Hooimeijer, and M. Ernst. HAMPI: A solver for string constraints. In *Proceedings of ISSTA*, 2009.
- [14] A. Krauss. *Automating recursive definitions and termination proofs in higher order logic*. PhD thesis, TU Munich, 2009.
- [15] K. R. Leino. Automating induction with an SMT solver. In *Proceedings of VMCAI*, volume 7148 of *LNCS*. Springer, 2012.
- [16] K. R. Leino and P. Ruegger. A polymorphic intermediate verification language: Design and logical encoding. In *Proceedings of TACAS*, volume 6015 of *LNCS*, 2010.
- [17] E. Meijer, M. Fokkinga, and R. Paterson. Functional programming with bananas, lenses, envelopes, and barbed wire. In *Proceedings of FPCA*, volume 523 of *LNCS*, 1991.
- [18] S. Miller, M. Whalen, and D. Cofer. Software model checking takes off. *CACM*, 53:58–64, February 2010.
- [19] R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML (Revised)*. The MIT Press, 1997.
- [20] M. Myreen. *Formal verification of machine-code programs*. PhD thesis, University of Cambridge, 2009.
- [21] S. Peyton Jones et al. The Haskell 98 language and libraries: The revised report. *Journal of Functional Programming*, 13(1):0–255, Jan 2003.
- [22] N. Schirmer. *Verification of sequential imperative programs in Isabelle/HOL*. PhD thesis, TU Munich, 2006.
- [23] P. Sestoft. ML pattern match compilation and partial evaluation. In *Dagstuhl Seminar on Partial Evaluation*, volume 1110 of *LNCS*, pages 446–464, 1996.
- [24] K. Slind and M. Norrish. A brief overview of HOL4. In

*Proceedings of TPHOLs*, volume 5170 of *LNCS*, pages 28–32, 2008.

- [25] P. Suter, M. Dotta, and V. Kuncak. Decision procedures for algebraic data types with abstractions. In *Proceedings of POPL*, pages 199–210. ACM, 2010.
- [26] P. Suter, A. Köksal, and V. Kuncak. Satisfiability modulo recursive programs. In E. Yahav, editor, *Proceedings of Static Analysis*, volume 6887 of *LNCS*, pages 298–315. Springer, 2011.
- [27] M. Whalen, D. Greve, and L. Wagner. Model checking information flow. In D. Hardin, editor, *Design and Verification of Microprocessor Systems for High-Assurance Applications*. Springer, 2010.