

Imperative Functional Programming with Isabelle/HOL

Lukas Bulwahn¹, Alexander Krauss¹, Florian Haftmann^{1*},
Levent Erkök², John Matthews²

¹ Technische Universität München,
Institut für Informatik, Boltzmannstraße 3, 85748 Garching, Germany

² Galois Inc., Beaverton, OR 97005, USA

Abstract. We introduce a lightweight approach for reasoning about programs involving imperative data structures using the proof assistant Isabelle/HOL. It is based on shallow embedding of programs, a polymorphic heap model using enumeration encodings and type classes, and a state-exception monad similar to known counterparts from Haskell. Existing proof automation tools are easily adapted to provide a verification environment. The framework immediately allows for correct code generation to ML and Haskell. Two case studies demonstrate our approach: An array-based checker for resolution proofs, and a more efficient bytecode verifier.

1 Introduction

A very common way of verifying programs in a HOL theorem prover is to use a shallow embedding and express the program as a set of recursive functions. Properties of the program can then be proved by induction. Despite some well-known limitations, shallow embeddings are widely used for verification. This success is due in part to the simplicity of the approach: A full-blown formal model of the operational or denotational semantics of the language is not required, and many technical difficulties (e.g. the representation of binders) are avoided altogether. Furthermore, the proof methods used are just standard induction principles and equational reasoning, and no specialized program logic is necessary. The specifications may be turned into executable code directly by means of code generation.

Until recently, this approach has been used primarily for purely functional programs. As the notion of side-effect is alien to the world of HOL functions, programs with imperative updates of references or arrays cannot be expressed directly. However, there are many examples where for efficiency's sake imperative data structures are unavoidable to obtain practically usable executable programs.

We aim to permit Haskell's imperative specification style in Isabelle/HOL [11], where local state references and mutable arrays can be dynamically allocated without having to add their types to the enclosing function's type signature [6]. From such specifications we then generate efficient imperative functional code. Currently we need to restrict the contents of references and mutable arrays to first order values, but this is still sufficient for many applications.

Accordingly, the contributions of this paper are:

* Supported by DFG project NI 491/10-1

1. A purely definitional polymorphic *heap* allowing to encode dynamic allocation of polymorphic first-order references and mutable arrays (§2).
2. A Haskell-style *heap monad* encapsulating the primitive heap operations and supporting abnormal termination through exceptions (§3); an adaption for Isabelle’s code generator allows to generate monadic Haskell and imperative ML³ code (§4).
3. A set of proof rules that allows to reason about such monadic programs (§5).
4. Two case studies (§6): an imperative MiniSat proof replay oracle and an imperative Jinja bytecode verifier.

1.1 Related Work

Since the seminal paper by Peyton Jones and Wadler [13], the use of monads to incorporate effects in purely functional programs is standard. However, up to now, no practically usable verification framework for such monadic programs exists.

For imperative programs, there are such tools: The Why/Krakatoa/Caduceus toolset [3] works by translating the source language into an intermediate language and using a verification condition generator to generate proof obligations. Schirmer [14] proposes a similar method, which is closely integrated with Isabelle/HOL, and whose metatheory is formally verified. These approaches rely on Hoare logic and a verification condition generator. The actual reasoning then happens on the generated verification conditions and is often outsourced to automatic provers. The user must provide enough annotations in the source code that the verification conditions can be solved by the automated prover. In our approach, reasoning happens on the source code level, which we find better suited for interactive use. Proof principles are similar to those used for purely functional programs, i.e. induction and equational reasoning.

Probably closest to our work is the concept of *single-threaded objects* [1] in the ACL2 prover. By declaring an object as single-threaded (and obeying rigorous syntactic restrictions), one instructs the prover to replace non-destructive updates by destructive ones. The rules ensure that referential transparency is not violated, and thus the code can be treated as purely functional in the reasoning phase. Our approach is similar in the sense that our heap can be seen as a single-threaded object. However, we allow the dynamic allocation of arrays and references, whereas in ACL2 imperative fields must be statically declared in a single record.

Imperative language features have previously been embedded in higher order logic via a *state monad* [8, 15]. As long as the monad primitives do not duplicate the state, the resulting programs are single threaded and can be safely transformed to monadic Haskell or imperative ML code. However, just like single-threaded objects, the state monad approach requires the state record be statically declared as part of the monad type itself, either fixed or as an explicit type parameter. This makes it difficult to write specifications that dynamically allocate new references or mutable arrays, or to compose monadic specifications that work over different state types.

Our heap model has some similarities to the one used by Tuch, Klein, and Norrish [16], especially concerning the use of type classes and phantom types to manage encodings. On the other hand, our model is slightly more abstract, since we are only dealing with functional languages instead of low level C code.

³ in its two flavors SML and OCaml

Nanevski et al. [10] describe how Hoare logic can be integrated in dependent type theory, yielding *Hoare Type Theory*, with a sophisticated type system and program logic. However, it seems that this requires significant modifications to the theorem prover in order to support such a system. In contrast, our approach was developed on top of standard Isabelle/HOL.

1.2 A trivial example: Array reversal

To illustrate imperative functional programming in monadic HOL, we define two simple functions, one for swapping two elements in an array, and one for reversing an array:

```

swap :: (α::hrep) array ⇒ nat ⇒ nat ⇒ unit Heap
swap a i j = do x ← a[i];
              y ← a[j];
              a[i] := y;
              a[j] := x;
              return ()

rev :: (α::hrep) array ⇒ nat ⇒ nat ⇒ unit Heap
rev a i j = (if i < j then do swap a i j;
              rev a (i + 1) (j - 1)
             else return ())

```

This idiom is well-known from Haskell: Manipulations of imperative arrays are monadic actions (of type α Heap), and they can be composed into more complex actions using the sequencing operation provided by the monad. Other language constructs (conditionals, recursion, data types) are taken from the functional part of the language.

Let us prove a lemma that describes the behavior of swap:

$$\begin{aligned}
&(h, h', r) \in \llbracket \text{swap } a \ i \ j \rrbracket \implies \\
&\text{get-array } a \ h' \ ! \ k = \\
&\quad (\text{if } k = i \ \text{then } \text{get-array } a \ h \ ! \ j \\
&\quad \ \ \text{else if } k = j \ \text{then } \text{get-array } a \ h \ ! \ i \ \text{else } \text{get-array } a \ h \ ! \ k)
\end{aligned}$$

The notation $\llbracket c \rrbracket$ stands for the big-step semantics of a command c , which is a ternary relation: $(h, h', r) \in \llbracket c \rrbracket$ holds iff the computation c started on heap h does not generate an exception and yields the result r and the updated heap h' . The lemma expresses how the entries of the updated array are related to original entries. The function *get-array* returns the contents of an array on a given heap as a list, and infix $!$ denotes list indexing.

The lemma is proved by unfolding the call to swap and applying standard rules for the semantics of the monad operations and the basic commands, which can easily be automated using existing Isabelle technology. Now let us turn to the reversal function:

$$\begin{aligned}
&(h, h', r) \in \llbracket \text{rev } a \ i \ j \rrbracket \implies \\
&\text{get-array } a \ h' \ ! \ k = \\
&\quad (\text{if } k < i \ \text{then } \text{get-array } a \ h \ ! \ k \\
&\quad \ \ \text{else if } j < k \ \text{then } \text{get-array } a \ h \ ! \ k \ \text{else } \text{get-array } a \ h \ ! \ (j - (k - i)))
\end{aligned}$$

Since *rev* is defined recursively, we proceed by induction. The proof is as one would expect: In the step case, we distinguish the cases $k < i$, $k = i$, $i < k < j$, $k = j$ and $j < k$, and apply the induction hypothesis and the lemma about swap.

1.3 Dynamic allocation: Linked lists

The ability to explicitly allocate memory is another fundamental technique in imperative programming. To illustrate how this idiom can be coded in HOL, we will show how to build and traverse a dynamically allocated linked list.

Linked lists are represented by a recursive datatype, where the tail of the list is a mutable reference.

```
datatype  $\alpha$  node = Empty | Node  $\alpha$  ( $\alpha$  node ref)
```

To convert a HOL list of elements to a linked list, we simply recurse over each tail, allocating the nodes as we go along by calling the *ref* function:

```
make-llist :: ( $\alpha$ ::hrep) list  $\Rightarrow$   $\alpha$  node Heap  
make-llist [] = return Empty  
make-llist (x:xs) = do tl  $\leftarrow$  make-llist xs;  
                    next  $\leftarrow$  ref tl;  
                    return (Node x next)
```

In the other direction, we can traverse a linked list as follows:

```
traverse :: ( $\alpha$ ::hrep) node  $\Rightarrow$   $\alpha$  list Heap  
traverse Empty = return []  
traverse (Node x r) = do tl  $\leftarrow$  !r;  
                    xs  $\leftarrow$  traverse tl;  
                    return (x:xs)
```

Note that the definitions of *make-llist* and *traverse* operationally mimic their equivalents in Haskell using the state monad, or in ML using imperative features.⁴

For reference, here is the *traverse* function as rendered by the code generator of our framework in ML:⁵

```
datatype 'a node = Node of 'a * 'a node ref | Empty;  
  
fun traverse A_ (Node (x, r)) =  
  (fn a as () =>  
    let  
      val tl = (fn () => ! r) ();  
      val xs = traverse A_ tl ();  
    in  
      x :: xs  
    end)  
| traverse A_ Empty = (fn () => []);
```

⁴ Technical details on the definition of *traverse* can be found in §7.3

⁵ The *A_* argument denotes the dictionary, which is not used in this particular example. See [4]

2 Modeling a polymorphic heap

In the following two sections we present our definitional model of a typed heap and the monad we are using. We present the theory in a bottom-up manner, and explain the most important design decisions.

Essentially, our heap will be a mapping $h :: \mathbb{N} \Rightarrow \mathcal{Val}$ from addresses to values. However, since values can generally have arbitrary types, this is difficult to model in a simply typed language. Since there is no HOL type that can contain all types, we are facing the problem what type to choose for \mathcal{Val} .

This problem could be solved using dependent types, but we want to stay in the simply typed framework, so we make a draconian restriction: We decide that functions cannot be stored on the heap, and use the natural numbers as value type, in which all first-order data objects will be encoded. We'll use phantom types (§2.2) to safely omit these encodings from the generated code.

Obviously this restrictive design decision excludes a fair number of relevant programs. But even then our model allows for interesting applications. Possibilities for lifting this restriction are discussed in §7.2.

2.1 Representable Types

Using encodings to circumvent restrictions in the type system seems very awkward at first, but we can make this transparent to the user by defining an axiomatic type class *countable*, with an axiom stating that the type can be encoded into the natural numbers:

$$\begin{aligned} \text{axclass } \textit{countable} &\subseteq \textit{type} \\ &\exists (\textit{enc} :: \alpha \Rightarrow \textit{nat}). \textit{inj } \textit{enc} \end{aligned}$$

Obviously, basic types like *nat*, *int* and all finite types are countable, and the well-known constructions can be used to show that if α and β are countable then so are $\alpha \times \beta$ and α *list*. In fact, such instance proofs are straightforward for first-order recursive data types and could be automated. The overloaded encoding and decoding functions are called *to-nat* and *from-nat*:

$$\begin{aligned} \textit{to-nat} &:: (\alpha :: \textit{countable}) \Rightarrow \textit{nat} \\ \textit{from-nat} &:: \textit{nat} \Rightarrow (\alpha :: \textit{countable}) \end{aligned}$$

2.2 Typed References

References are just addresses, i.e. natural numbers

$$\text{datatype } \alpha \textit{ ref} = \textit{Ref } \textit{nat}$$

with the projection $\textit{addr-of } (\textit{Ref } n) = n$. Here, unlike above, the type system is again a useful tool instead of just a handicap: The *phantom type* α , which does not occur on the right hand side of the definition, allows us to view references as typed objects as we know them from ML, although the underlying representation is untyped.

Reference equality We will certainly need to reason about reference (in)equality. For example, we would expect the following simplification rule to hold, where r and s are references and h is a heap:

$$r \neq s \implies \text{get-ref } r \ (\text{set-ref } s \ x \ h) = \text{get-ref } r \ h$$

Indeed, when we write down and prove this lemma, everything seems to work. Only if we look at the inferred types, there is an unpleasant surprise: Because of the equality in the premise, the references r and s have the same type, and we have thus only proved a special case. Of course we want to perform the same simplification if we have references of different types, and ideally, we want the condition $r \neq s$ to be immediate, when the references have different types.

The solution is to define a heterogeneous inequality relation for references, which just strips away the phantom type and compares the bare addresses:

$$r \not\approx s \leftrightarrow (\text{addr-of } r \neq \text{addr-of } s) \text{ }^6$$

If r and s have the same type, the relation $\not\approx$ coincides with \neq .

2.3 Type reflection

Comparing references of different types is a little artificial. In a typed language, aliasing between e.g. an integer and a boolean reference is not possible, and the above rewrite rule should be applicable unconditionally. Our model will be built in such a way that we can automatically derive $r \not\approx s$, whenever r and s have different types:

We define a type *typerep* and a type class *typeable* to reflect the syntax of (monomorphic) types back into the language of terms:

datatype *typerep* = *Typerep string (typerep list)*

class *typeable* = *type* +
fixes *typerep* :: $\alpha \text{ itself} \Rightarrow \text{typerep}$

The predefined type $\alpha \text{ itself}$ comes with a singleton term written $TYPE(\alpha)$ which is used to embed types into terms. We write $RTYPE(\alpha)$ for *typerep* ($TYPE(\alpha)$). The overloaded function *typerep* constructs a concrete syntactic representation of a type name. Its definition for concrete types is completely schematic (and easily automated):

$$\begin{aligned} RTYPE(\text{nat}) &= \text{Typerep } \text{"nat"} \ [] \\ RTYPE(\text{bool}) &= \text{Typerep } \text{"bool"} \ [] \\ RTYPE(\alpha \ \text{list}) &= \text{Typerep } \text{"list"} \ [RTYPE(\alpha)] \end{aligned}$$

The result of this exercise (which is also common in the Haskell world) is that we can now compare types for equality explicitly. For example: $RTYPE(\text{nat}) \neq RTYPE(\text{bool})$ and $RTYPE(\text{char list}) = RTYPE(\text{string})$ are theorems⁷, however $RTYPE(\alpha) \neq RTYPE(\beta)$ is not, since α and β could later be instantiated to the same type.

Now we can refine the definition of reference inequality as follows:

⁶ The \leftrightarrow denotes equality of *bool* values

⁷ Note that in Isabelle, *string* just abbreviates *char list* on the surface syntax level

$$(r :: \alpha \text{ ref}) \not\cong (s :: \beta \text{ ref}) \leftrightarrow \text{RTYPE}(\alpha) \neq \text{RTYPE}(\beta) \vee \text{addr-of } r \neq \text{addr-of } s$$

From this immediately follows that references of different types are always unequal. Hence, aliasing proof obligations like $r \not\cong s$ can be solved automatically, if r and s are of different types.

2.4 The Heap

The heap is modelled as a mapping from type representations and addresses to *nat*-encoded values. We use two separate mappings for references and arrays (which are mapped to lists of encoded values). Additionally, we use a counter which bounds the currently used address space. It is incremented when new references are created.

```
record heap =
  refs :: typerep  $\Rightarrow$  addr  $\Rightarrow$  nat
  arrays :: typerep  $\Rightarrow$  addr  $\Rightarrow$  nat list
  lim :: addr
```

We can now define the basic heap operations, such as allocation, reading and writing of references. Note how the embeddings and projections are used here to convert the stored values into their respective encodings and back⁸:

```
get-ref :: ( $\alpha$ ::hrep) ref  $\Rightarrow$  heap  $\Rightarrow$   $\alpha$ 
get-ref (Ref r) (h(|refs := f|)) = from-nat (f RTYPE( $\alpha$ ) r)

set-ref :: ( $\alpha$ ::hrep) ref  $\Rightarrow$   $\alpha$   $\Rightarrow$  heap  $\Rightarrow$  heap
set-ref (Ref r) x (h(|refs := f|)) =
  h(|refs := f(RTYPE( $\alpha$ ) := (f RTYPE( $\alpha$ ))(r := to-nat x))|)

new-ref :: heap  $\Rightarrow$  ( $\alpha$ ::hrep) ref  $\times$  heap
new-ref (h(|lim := l|)) = (Ref l, h(|lim := Suc l|))
```

Operations for arrays are analogous. From these definitions, we can now easily prove the expected lemmas, expressing the interaction of the operations, e.g.:

```
get-ref r (set-ref r x h) = x
r  $\not\cong$  s  $\implies$  get-ref r (set-ref s x h) = get-ref r h
```

Since arrays and references occupy different heap areas, the corresponding heap operations always commute:

```
get-array a (set-ref r x h) = get-array a h
get-ref r (set-array a xs h) = get-ref r h
```

⁸ The *hrep* type class just intersects the classes *countable* and *typeable*; also note that the $h(|\dots := \dots|)$ syntax denotes component assignment on records and may also be used for pattern matching.

3 The Heap Monad

We now define a monad which characterizes computations affecting the heap. An imperative program with return type α will be logically represented as a value of type α *Heap*. Essentially, our monad is a state-exception monad, where the state is the heap from the previous section:

```
datatype  $\alpha$  Heap = Heap (heap  $\Rightarrow$  ( $\alpha$  + exception)  $\times$  heap)
heap f      = Heap ( $\lambda h. \mathbf{let}$  (x, h') = f h in (Inl x, h'))
execute (Heap f) = f
```

Exceptions are essentially strings generated by *error* :: *string* \Rightarrow *exception* and are not caught inside the monad; they are a mere device to introduce a notion of abnormal termination. The monad operations *return*, ($\gg=$) and *raise* are defined as expected:

```
return x = heap ( $\lambda h. (x, h)$ )
f  $\gg=$  g = Heap
          ( $\lambda h. \mathbf{case}$  execute f h of (Inl x, h')  $\Rightarrow$  execute (g x) h'
          | (Inr e, h')  $\Rightarrow$  (Inr e, h'))
raise s = Heap ( $\lambda h. (\mathbf{Inr}$  (error s), h)
```

Isabelle's syntax facilities allow for Haskell-style do-notation. Lifting the heap operations into the monad is straightforward:


```
ref x   = heap ( $\lambda h. \mathbf{let}$  (r, h') = new-ref h in (r, set-ref r x h'))
!r      = heap ( $\lambda h. (\mathbf{get-ref}$  r h, h)
r := x = heap ( $\lambda h. ((\mathbf{)}, \mathbf{set-ref}$  r x h))

array n x = heap ( $\lambda h. \mathbf{let}$  (a, h') = new-array h in (a, set-array a (replicate n x) h'))
length a = heap ( $\lambda h. (\mathbf{|get-array}$  a h, h)
a[i]     = do len  $\leftarrow$  length a;
              (if i < len then heap ( $\lambda h. (\mathbf{get-array}$  a h ! i, h)
              else raise "array lookup: index out of range")
a[i] := x = do len  $\leftarrow$  length a;
              (if i < len then heap ( $\lambda h. (\mathbf{a, set-array}$  a (get-array a h[i := x]) h)
              else raise "array update: index out of range")
```

These are the necessary foundations to write stateful programs like in §1.2.

4 Execution

When we consider some parts of HOL as the shallow embedding of a programming language, then the inverse of that embedding is called *code generation*. Isabelle's code generator [4] can produce SML, OCaml and Haskell code from executable HOL specifications. In a first approximation, the executable fragment of HOL consists of datatype and function definitions, which are simply translated to their counterparts. This guarantees partial correctness: if $\langle s \rangle$ denotes the generated code from term *s*, then each abstract



HOL	ML	Haskell
α Heap	$unit \Rightarrow \alpha$	$ST \xi \alpha$
$t \gg= (\lambda x. f)$	$\lambda(). let x = t () in f () end$	$t \gg= (\lambda x. f)$
return t	$\lambda(). t$	return t
α ref	α ref	$STRef \xi \alpha$
ref x	$\lambda(). ref x$	newSTRef x
! r	$\lambda(). !r$	readSTRef r
$r := t$	$\lambda(). r := x$	writeSTRef $r x$
α array	α array	$STArray \xi Integer \alpha$
array $n x$	$\lambda(). Array.array (n, x)$	newArray $(0, n) x$
$a[i]$	$\lambda(). Array.sub (a, i)$	readArray $a i$
$a[i] := x$	$\lambda(). Array.update (a, i, x)$	writeArray $a i x$
length a	$\lambda(). Array.length a$	liftM snd (getBounds a)
raise s	$\lambda(). raise Fail s$	error s

(b) Translating monadic constructs to ML and Haskell

Fig. 1. Code generation

evaluation step from $\langle s \rangle$ to some t' in the target language corresponds to an equational rewrite step $s = t$ in HOL, such that $\langle t \rangle = t'$ (cf. Fig. 1(a)).

The reference and array operations are mapped to the target language as given in Fig. 1(b). Since ML expressions may already contain side effects, the monad vanishes and is just replaced by a unit abstraction to ensure the correct evaluation order.

For Haskell we use the built-in ST state monad, together with the corresponding $STRef$ and $STArray$ types. Recall that our HOL programs only raise exceptions but never handle them – instead of dealing with them inside the monad, we treat them as partiality, using the $error$ primitive.

Note that the extended executable fragment of HOL does not include the constructions that were used to define the heap monad: If we break the monad abstraction (e.g. by writing $heap (\lambda h. (h, h))$), the results are no longer executable and trying to generate code for them causes an error, just like for other purely logical notions like quantifiers.

5 Verification

Having defined the model of execution for our stateful programs, we need verification tools which can be used to prove an individual program correct. Our model does not force us to use a particular technique: We can choose any calculus (like e.g. Hoare logic) that is sound with respect to the semantics we have defined. After a bit of experimenting, we opted for a very simple method, which seems to fit well with the structured proof language Isabelle/Isar. A deeper comparison of this relational style with Hoare logic is beyond the scope of this paper.

We use the relational description of the big-step semantics we have already seen in §1.2. The relation is defined by:

$$(h, h', r) \in \llbracket c \rrbracket = ((\text{Inl } r, h') = \text{execute } c \ h)$$

We can prove rules which connect this relation to the different basic commands. Here is the rule for the bind operation.

$$\begin{aligned} (h, h'', r') \in \llbracket f \gg= g \rrbracket &\implies \\ (\bigwedge h' r. (h, h', r) \in \llbracket f \rrbracket \implies (h', h'', r') \in \llbracket g \ r \rrbracket \implies P) &\implies P \end{aligned}$$

Note the elimination rule format. Since the $(\dots) \in \llbracket \dots \rrbracket$ relation usually lives in the premise of a statement, we use elimination rules to manipulate it: if our goal has a premise of the form $(h, h'', r) \in \llbracket f \gg= g \rrbracket$, we can obtain the intermediate heap h' and the new assumptions $(h, h', a) \in \llbracket f \rrbracket$ and $(h', h'', r) \in \llbracket g \ a \rrbracket$. These elimination rules allow us to systematically decompose compound statements into primitive steps. Here are some other rules:

$$\begin{aligned} (h, h', r) \in \llbracket \text{return } x \rrbracket &\implies (r = x \implies h = h' \implies P) \implies P \\ (h, h', r) \in \llbracket a[i] \rrbracket &\implies \\ (r = \text{get-array } a \ h \ ! \ i \implies h = h' \implies i < \text{length-array } a \ h \implies P) &\implies P \\ (h, h', r) \in \llbracket a[i] := v \rrbracket &\implies (r = a \implies h' = \text{Heap.upd } a \ i \ v \ h \implies P) \implies P \end{aligned}$$

By feeding these rules into Isabelle's *auto* method, we obtain a reasonable *ad-hoc* automation, which makes proofs quite short.

6 Case studies

6.1 A SAT Checker

Our first case study is motivated by the wish to integrate SAT solvers into Isabelle in a scalable way, such that they can be used to solve large propositional proof obligations.

We aim at a compromise between performing a full replay of the proof within Isabelle and trusting the SAT solver completely. The first approach was taken by Weber and Amjad [17] and gives the usual high assurance of the LCF principle, but is computationally expensive. On the other end of the spectrum, trusting the external tool is obviously cheap but unsatisfactory.

A reasonable compromise is to run the external proof (a standard propositional resolution proof) through a checker, which is itself proved correct in Isabelle. This gives a good balance between assurance and cost, since unlike the SAT solver, the checker is formally verified, and checking a proof is about an order of magnitude faster than replaying the inferences in Isabelle.

Usually, for such a reflective approach, the checker would need to be purely functional. Using our framework, we can implement a checker that uses destructive arrays instead, which gives us another 30% speedup over a purely functional implementation with balanced trees.

The core of our checker operates on a table that stores the clauses that have already been derived. Clauses are modeled (purely functionally) as sorted lists of integers, where a negative number signifies a negated variable:

types		datatype <i>ProofStep</i> =
<i>idx</i>	= <i>nat</i>	<i>Root idx clause</i>
<i>lit</i>	= <i>int</i>	<i>Resolve idx resolvants</i>
<i>clause</i>	= <i>lit list</i>	<i>Delete idx</i>
<i>resolvants</i>	= <i>idx × (lit × idx) list</i>	

A proof step can either (a) add a new so-called *root clause* to the array, (b) derive a new clause from existing clauses and store it in the array, or (c) delete a clause from the array, to free some memory.

The root clauses are the initial clauses from which a contradiction is derived. It is a specialty of the MiniSAT [2] proof format that root clauses may be added any time during the proof, hence our checker must accumulate all root clauses it encounters in a list. Then, if the checker succeeds in deriving the empty clause, the root clauses it has collected must be inconsistent.

A *Resolve* step derives a new clause in a series of resolutions: *Resolve i (j, rs)* starts with clause no. *j* and resolves it with the clause/variable pairs in *rs*. In the end, the result is stored at position *i*. The *Delete* proof step removes a clause from the array. This weakening step is simply an optimization to reduce memory usage of the checker by removing clauses that are no longer needed.

With clauses modeled as sorted lists, resolution is essentially a merge operation and can be done in just one traversal. However, the operation may fail if the literal does not occur in the clause. It is convenient to let the monad deal with such failures, even if no heap access is required. Hence our *resolve* operation has the following type (for brevity, we omit the implementation, which does not contain surprises):

$$\text{resolve} :: \text{lit} \Rightarrow \text{clause} \Rightarrow \text{clause} \Rightarrow \text{clause Heap}$$

The function *get-clause* retrieves a clause from the array. It fails if it sees a *None*:

$$\text{get-clause} :: \text{clause option array} \Rightarrow \text{idx} \Rightarrow \text{clause Heap}$$

The heart of our checker is the function *step*, which processes a single proof step, collecting root clauses in the accumulator list *rCs*:

$$\text{step} :: \text{clause option array} \Rightarrow \text{ProofStep} \Rightarrow \text{clause list} \Rightarrow \text{clause list Heap}$$

$$\text{step } a \text{ (Root } cid \text{ clause) } rCs = \text{do } a[cid] := \text{Some (remdups (sort clause));}$$

$$\text{return (clause.rCs)}$$

$$\text{step } a \text{ (Resolve saveTo (i, rs)) } rCs =$$

$$\text{do } cli \leftarrow \text{get-clause } a \text{ i;}$$

$$\text{result} \leftarrow \text{foldM } (\lambda(l, j) \text{ c. get-clause } a \text{ j} \gg= \text{resolve } l \text{ c}) \text{ rs } cli;$$

$$a[\text{saveTo}] := \text{Some result;}$$

$$\text{return } rCs$$

$$\text{step } a \text{ (Delete } cid \text{) } rCs = \text{do } a[cid] := \text{None;}$$

$$\text{return } rCs$$

Finally, a wrapper function *checker* just allocates an array of a given size, folds the step function over a list of proof steps, and finally checks for the empty clause at some given position. Our main result is the following partial correctness theorem:

$$(h, h', cs) \in \llbracket checker\ n\ p\ i \rrbracket \implies inconsistent\ cs$$

Integration Since we have verified our checker, we may now choose to use it to import proofs into Isabelle. This can be done using a generic *monadic evaluation oracle*, which implements the following inference rule:

$$\frac{\bigwedge h\ h'. (h, h', r) \in \llbracket c \rrbracket \implies P\ r}{P\ r} \text{ (if } c, \text{ when executed in ML, evaluates to } r)$$

Thus we can discharge the premise of a partial correctness theorem by just running the generated code in ML.

The soundness of this rule relies on the assumption that the semantics of ML is compatible with our model of monadic programs. At the moment, we have no proof of this assumption.

However, such a generic reflection mechanism, which provides a clearly defined way to extend the theorem prover by reflected imperative proof tools, still provides higher assurance than an ad-hoc extension, since the monadic code is verified, and no additional “glue code” is required for the integration.

In particular, nothing in our particular development of the SAT checker needs to be trusted.

6.2 A Jinja Bytecode verifier

Our second case study is a modification of the Jinja bytecode verifier. Jinja [7] is a complete formal model of a Java-like language, which includes a formal semantics, type system, virtual machine model, compiler, and bytecode verifier.

Essentially, the bytecode verifier performs an abstract interpretation of the bytecode instructions, keeping track of the abstract state, that is, the types of values in registers and on the stack. The central data structure is a mapping that assigns such an abstract state to every bytecode instruction. Then, this information is propagated to the successors of the instruction until a fixed point is reached.

In the existing implementation, this mapping is represented by a list of fixed length. In our modification, we use an imperative array instead, with the obvious advantages: constant-time access and no garbage.

Fortunately, the bytecode verifier is modeled in a very abstract framework using a semilattice (type σ), which hides all the technical details of the virtual machine. Later, the “real thing” can be obtained by instantiation. A bytecode method is modelled by a function $step :: nat \Rightarrow \sigma \Rightarrow (nat \times \sigma) list$ that maps a given program position and an abstract machine state to a list of possible successor positions and states. Additional requirements for the *step* function (e.g. monotonicity) are detailed in [7].

Figure 2 shows the pure version of the bytecode verifier together with its monadic counterpart. This side-by-side comparison shows that the differences between the two versions are small. Consequently, proving partial correctness of *kildallM* wrt. *kildall* is straightforward:

$$\tau s \in \text{list } n A \implies (h, h', \tau s') \in \llbracket \text{kildallM } \tau s \rrbracket \implies \tau s' = \text{kildall } \tau s$$

This shows that it is relatively easy to move from a purely functional specification to a monadic one, which can then be executed efficiently.

<pre> propa [] $\tau s w = (\tau s, w)$ propa ((q, τ)·qs) $\tau s w =$ ($\text{let } u = \tau \sqcup \tau s ! q;$ $w' = \text{if } u = \tau s ! q \text{ then } w \text{ else } \{q\} \cup w$ $\text{in propa } qs (\tau s[q := u]) w')$ </pre>	<pre> propaM [] $\tau s w = \text{return } w$ propaM ((q, τ)·qs) $\tau s w =$ ($\text{do } \tau' \leftarrow \tau s[q];$ $\text{let } u = \tau \sqcup \tau';$ $\text{let } w' = (\text{if } u = \tau' \text{ then } w \text{ else } \{q\} \cup w);$ $\tau s[q] := u;$ $\text{propaM } qs \tau s w')$ </pre>
<pre> iter ($\tau s, w$) = ($\text{if } w = \emptyset \text{ then } \tau s$ $\text{else let } p = \text{SOME } p. p \in w$ in iter $(\text{propa } (\text{step } p (\tau s ! p)) \tau s (w - \{p\})))$ </pre>	<pre> iterM $\tau s w =$ ($\text{if } w = \emptyset \text{ then freeze } \tau s$ $\text{else let } p = \text{SOME } p. p \in w$ $\text{in do } v \leftarrow \tau s[p];$ $w' \leftarrow \text{propaM } (\text{step } p v) \tau s$ $(w - \{p\});$ $\text{iterM } \tau s w')$ </pre>
<pre> kildall $\tau s = \text{iter } (\tau s, \text{unstabes } \tau s)$ </pre>	<pre> kildallM $\tau s =$ ($\text{do } a \leftarrow \text{of-list } \tau s;$ $\text{iterM } a (\text{unstabes } \tau s)$ </pre>

Fig. 2. Pure vs. monadic versions of the bytecode verifier

7 Problems and Limitations

7.1 No Monad Polymorphism

Of course, one would like to specify a monad as a constructor class, and see our heap monad just as a particular instance of the general concept. However, for this we would need type constructor polymorphism, which is not supported in HOL. We must be satisfied with the possibility of defining concrete instances of monads.

Huffman, Matthews, and White [5] describe how to simulate constructor classes in an extension of HOL, but their embedding does not seem practical for our application.

7.2 Heap model

Our simple heap model prohibits storing any kind of function value in mutable references. Although many applications can live with this limitation, it may be painful in other situations. One can think of different ways to improve this situation:

Encoding types of order n . Just like we now encode all first-order values in \mathbb{N} , one could also encode all functions on such values by $\mathbb{N} \Rightarrow \mathbb{N}$, and all functions that take

such arguments by $(\mathbb{N} \Rightarrow \mathbb{N}) \Rightarrow \mathbb{N}$, and so on. For any given order, we can encode all “smaller” types in a single type. Again, this can be made transparent using type classes. Probably, order 3 or 4 would be enough for most practical purposes.

Dependent types In a dependently typed system, one could do without explicit encodings, and represent heap values as a dependent pair of a type and a value. In such a system, the type *heap* would live in some higher universe than the types used in a program.

ZF extension HOLZF [12] is a consistent extension of HOL which declares a set-theoretic universe Z , in which all HOL types can be embedded. In such a system we could store the full tower of (pure, monomorphic) higher order functions over the naturals, since our heap function could take values in Z .

However, even these extensions will not allow us to store *monadic* functions in the heap. The collection of heap monad functions has at least the cardinality of $heap \Rightarrow heap$, which is strictly larger than *heap* itself in classical HOL.

One avenue of escape would be to limit ourselves to the constructive portion of HOL and build some kind of impredicative datatype facility to represent the heap. A more pragmatic option is to store only a representable subset of the full function space in the heap, for example just the continuous functions as is done in Isabelle/HOLCF[9]. We would retain the full power of classical HOL while still allowing to store all (partially) executable functions, which are the only ones we are really interested in.

7.3 Recursive Functions

Monadic functions can be defined recursively just like any other function by using the available packages in Isabelle. However, proving termination of the functions can sometimes be tricky, as the following example demonstrates:

$$\begin{aligned}
 f &:: \text{nat ref} \Rightarrow \text{nat} \Rightarrow \text{nat Heap} \\
 f \ r \ n &= \text{do } x \leftarrow !r; \\
 &\quad (\text{if } x = 0 \text{ then return } n \text{ else do } r := x - 1; \\
 &\quad \quad f \ r \ (\text{Suc } n))
 \end{aligned}$$

Since there is no wellfounded order for which $(r, \text{Suc } n) \prec (r, n)$ holds, we cannot hope to define f by wellfounded recursion on its arguments. Instead, the recursion happens on the heap itself, which is not an explicit argument of the function. To define f , we must first break the monad abstraction and define a function $f' :: \text{nat ref} \Rightarrow \text{nat} \Rightarrow \text{heap} \Rightarrow (\text{nat} + \text{exception}) \times \text{heap}$, which explicitly recurses over the heap. Then f can be defined in terms of f' , deriving the above recursion equation.

Another issue is that when building pointer structures on the heap, many functions are actually partial, since the structures can become cyclic. The attentive reader may have noticed that the *traverse* function in §1.3 is in fact such an example.

However, it turns out that even nonterminating recursive functions are definable if the recursion happens within the heap monad, since such definitions always have a total model. This argument is similar to the observation that tail-recursive functions can always be defined in HOL (e.g. using a *while* combinator). The details are beyond the scope of this paper, so we just mention that *traverse* was defined using a monadic recursion combinator *MREC*, which satisfies the following recursion equation:

$$\begin{aligned}
& \mathit{MREC} :: (\alpha \Rightarrow (\beta + \alpha) \mathit{Heap}) \Rightarrow (\alpha \Rightarrow \alpha \Rightarrow \beta \Rightarrow \beta \mathit{Heap}) \Rightarrow \alpha \Rightarrow \beta \mathit{Heap} \\
& \mathit{MREC} f g x = \\
& \mathit{do} y \leftarrow f x; \\
& \quad (\mathit{case} y \mathit{of} \mathit{Inl} r \Rightarrow \mathit{return} r \mid \mathit{Inr} s \Rightarrow \mathit{do} z \leftarrow \mathit{MREC} f g s; \\
& \quad \quad \quad g x s z)
\end{aligned}$$

In the future, we plan to provide automation for defining such recursive functions.

7.4 External I/O

Another practical limitation is that our heap monad does not support any kind of interaction with the outside world. This means, for example, that the full sequence of MiniSat proof steps needs to be passed into our SAT checker from the start. This becomes a problem for long-running proofs where the number of steps may exceed the total size of Isabelle’s memory.

However, if our monad supported IO actions then we could incrementally ask MiniSat to supply us just the next portion of the proof to check, and never have to represent the entire proof at once. Supporting I/O would require us to extend our heap model to include relevant aspects of the outside system, plus some kind of nondeterminism for I/O actions, to take into account that we can never model the world in its entirety.

8 Conclusion

We presented a lightweight approach to reuse our favorite theorem prover for verifying monadic programs that manipulate a state. Our shallow embedding of imperative constructs in HOL is a continuation of the traditional way of modeling programs and systems by recursive functions, which can be translated to “real” programs by a code generator. Although there are still some limitations (see §7), our case studies show that it is already quite useful in its current form. Equipped with that, we want to tackle specification, verification and prototypic code generation for compute-intensive applications like e.g. microprocessor models. Another important application is the extension of the Isabelle system itself by means of verified monadic proof procedures as we have sketched it in §6.1.

Future work will also focus on alleviating current limitations, most notably to allow a broader range of heap-representable types, monadic I/O, and more automation for defining monadic recursive functions.

9 Acknowledgments

We would like to thank David Hardin, Joe Hurd, Matt Kaufmann, Dylan McNamee, Tobias Nipkow, Konrad Slind, and Tjark Weber for their useful discussions, encouragement and feedback on our work.

References

- [1] R. S. Boyer and J. S. Moore. Single-threaded objects in ACL2. In *PADL '02: Proceedings of the 4th International Symposium on Practical Aspects of Declarative Languages*, pages 9–27, London, UK, 2002. Springer.
- [2] N. Een and N. Sörensson. An extensible sat-solver. In G. Goos, J. Hartmanis, and J. van Leeuwen, editors, *Theory and Applications of Satisfiability Testing*, volume 2919 of *LNCS*, pages 502 – 518. Springer, 2003.
- [3] J.-C. Filliâtre and C. Marché. The Why/Krakatoa/Caduceus platform for deductive program verification. In W. Damm and H. Hermanns, editors, *19th International Conference on Computer Aided Verification*, LNCS, Berlin, Germany, July 2007. Springer.
- [4] F. Haftmann and T. Nipkow. A code generator framework for Isabelle/HOL. Technical Report 364/07, Department of Computer Science, University of Kaiserslautern, 08 2007.
- [5] B. Huffman, J. Matthews, and P. White. Axiomatic constructor classes in Isabelle/HOLCF. In J. Hurd and T. F. Melham, editors, *TPHOLS*, volume 3603 of *LNCS*, pages 147–162. Springer, 2005.
- [6] S. P. Jones and J. Launchbury. Lazy functional state threads. In *SIGPLAN Conference on Programming Language Design and Implementation*, pages 24–35, 1994.
- [7] G. Klein and T. Nipkow. A machine-checked model for a Java-like language, virtual machine and compiler. *ACM Trans. Program. Lang. Syst.*, 28(4):619–695, 2006.
- [8] S. Krstić and J. Matthews. Verifying BDD algorithms through monadic interpretation. In *VMCAI '02: Revised Papers from the Third International Workshop on Verification, Model Checking, and Abstract Interpretation*, pages 182–195, London, UK, 2002. Springer.
- [9] O. Müller, T. Nipkow, D. v. Oheimb, and O. Slotosch. HOLCF = HOL + LCF. *Journal of Functional Programming*, 9:191–223, 1999.
- [10] A. Nanevski, G. Morrisett, and L. Birkedal. Polymorphism and separation in hoare type theory. In *ICFP '06: Proceedings of the eleventh ACM SIGPLAN international conference on Functional programming*, pages 62–73, New York, NY, USA, 2006. ACM.
- [11] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL – A Proof Assistant for Higher-Order Logic*. LNCS 2283. Springer, 2002.
- [12] S. Obua. Partizan games in Isabelle/HOLZF. In K. Barkaoui, A. Cavalcanti, and A. Cerone, editors, *ICTAC*, volume 4281 of *LNCS*, pages 272–286. Springer, 2006.
- [13] S. Peyton Jones and P. Wadler. Imperative functional programming. In *Proc. 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'93)*, pages 71–84, 1993.
- [14] N. Schirmer. A verification environment for sequential imperative programs in Isabelle/HOL. In F. Baader and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning*, volume 3452, pages 398–414. Springer, 2005.
- [15] C. Sprenger and D. A. Basin. A monad-based modeling and verification toolbox with application to security protocols. In K. Schneider and J. Brandt, editors, *TPHOLS*, volume 4732 of *LNCS*, pages 302–318. Springer, 2007.
- [16] H. Tuch, G. Klein, and M. Norrish. Types, bytes, and separation logic. In M. Hofmann and M. Felleisen, editors, *Proc. 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'07)*, pages 97–108, Nice, France, Jan. 2007.
- [17] T. Weber and H. Amjad. Efficiently checking propositional refutations in HOL theorem provers. *Journal of Applied Logic*, 2007. To appear.