

Recursive Definitions of Monadic Functions

Alexander Krauss

Technische Universität München, Institut für Informatik
<http://www.in.tum.de/~krauss>

Abstract

Using standard domain-theoretic fixed-points, we present an approach for defining recursive functions that are formulated in monadic style. The method works both in the simple option monad and the state-exception monad of Isabelle/HOL's imperative programming extension, which results in a convenient definition principle for imperative programs, which were previously hard to define.

For such monadic functions, the recursion equation can always be derived without preconditions, even if the function is partial. The construction is easy to automate, and convenient induction principles can be derived automatically.

1 Introduction

Tool support for non-primitive recursion in interactive theorem provers has made good progress in the last years. Although the base logic of most proof assistants has no support for general recursion or partial functions, tools exist to reduce such definitions to more basic principles in an automated and mostly transparent way [23, 2, 17, 5].

This paper discusses a class of definitions which are not yet supported by any existing tool: imperative computations wrapped up in a state monad. These functions present a challenge, since the actual structure of the recursion is not directly visible from the definition itself, but depends on the implicit state argument. Before considering imperative programs, we first develop our approach in the simpler option monad. The method can then be extended to state monads without much difficulty.

1.1 Notation

We work in the setting of Isabelle/HOL [20], which implements a variant of classical higher-order logic extended with type classes and overloading. Its syntax mostly conforms to standard mathematical notation, except for a few idiosyncracies that arise from the generic nature of the Isabelle system, notably the two versions of implication (\longrightarrow and \implies) and universal quantification (\forall and \bigwedge), corresponding to the meta and the object level. The reader can safely treat them as interchangeable for the purpose of this paper. Function types are written using \Rightarrow , and other basic types include *bool*, *nat* and α *list*.

1.2 Function Definitions in Isabelle/HOL

To explain the problem, we briefly outline the state of the art concerning function definitions in Isabelle/HOL [20]. The definition facilities (commonly called the *function package* [17, 16]) work as follows:

From the specification of a partial function given by a recursive equation $f\ x = F\ f\ x$, the function package produces a total function $f :: \sigma \Rightarrow \tau$, together with a domain predicate $dom :: \sigma \Rightarrow bool$, which models the set of values x where the recursion terminates. For values outside the domain, we can still write the term $f\ x$, but its value may be unknown or meaningless.

The original recursive specification is then derived as a theorem, constrained by the domain condition:

$$\text{dom } x \implies f x = F f x$$

This condition can be removed by proving that the function is total, i.e., $\forall x. \text{dom } x$. An induction rule is also derived, which is guarded by the domain condition as well.

The use of a predicate to describe terminating inputs is shared by various related approaches [9, 10, 6, 12]. In our classical simply-typed setting, it has the advantage that the function can be used syntactically as if it were total. In particular, one can treat a total function as partial temporarily, until its termination proof is finished and the predicate can be discarded. This has proved very useful for nested recursion. However, for truly partial functions the domain condition does not go away and has to be dealt with in proofs.

1.3 Imperative Functional Programs in Isabelle/HOL

Imperative HOL [7] is an extension of Isabelle/HOL which allows modeling and reasoning about programs that manipulate a heap. It defines a type *heap*, which models a store where references can be allocated and updated. (The details behind the type *heap* are omitted and can be safely ignored.)

$$\begin{aligned} \text{new-ref} &:: \text{heap} \Rightarrow \alpha \text{ ref} \times \text{heap} \\ \text{get-ref} &:: \alpha \text{ ref} \Rightarrow \text{heap} \Rightarrow \alpha \\ \text{set-ref} &:: \alpha \text{ ref} \Rightarrow \alpha \Rightarrow \text{heap} \Rightarrow \text{heap} \end{aligned}$$

Heap-modifying programs are modelled as monadic computations in the so-called *heap monad*:

$$\begin{aligned} \text{datatype } \alpha \text{ Heap} &= \text{Heap} (\text{heap} \Rightarrow (\alpha + \text{exception}) \times \text{heap}) \\ \text{return} &:: \alpha \Rightarrow \alpha \text{ Heap} \\ \text{return } x &= \text{Heap} (\text{Pair} (\text{Inl } x)) \\ \gg= &:: \alpha \text{ Heap} \Rightarrow (\alpha \Rightarrow \beta \text{ Heap}) \Rightarrow \beta \text{ Heap} \\ f \gg= g &= \text{Heap} (\lambda h. \text{case } \text{exec } f h \text{ of } (\text{Inl } x, h') \Rightarrow \text{exec } (g x) h' \mid (\text{Inr } \text{exn}, h') \Rightarrow \\ &(\text{Inr } \text{exn}, h')) \end{aligned}$$

where $\text{exec } (\text{Heap } f) = f$. This is nothing more than a state-exception monad, whose state type is *heap*. The singleton type *exception* is a simplistic way of modelling irrecoverable failure of the computation. The primitive heap operations are straightforwardly lifted to monadic operations with the following types:

$$\begin{aligned} \text{Ref.new} &:: \alpha \Rightarrow \alpha \text{ ref Heap} \\ \text{Ref.lookup} &:: \alpha \text{ ref} \Rightarrow \alpha \text{ Heap} \\ \text{Ref.update} &:: \alpha \text{ ref} \Rightarrow \alpha \Rightarrow \text{unit Heap} \end{aligned}$$

We abbreviate $\text{Ref.lookup } r$ by $!r$ and $\text{Ref.update } r x$ by $r := x$. Moreover, we use a notation similar to Haskell, i.e., $\text{do } x \leftarrow f; g x \text{ done}$ abbreviates $f \gg= (\lambda x. g x)$. For example, here is a data type of heap-allocated linked lists, and a function $\text{traverse} :: \alpha \text{ node} \Rightarrow \alpha \text{ list Heap}$ that traverses a linked list and turns it into an ordinary list:

$$\begin{aligned} \text{datatype } \alpha \text{ node} &= \text{Empty} \mid \text{Node } \alpha (\alpha \text{ node ref}) \\ \text{traverse } \text{Empty} &= \text{return } [] \\ \text{traverse } (\text{Node } x r) &= \text{do } tl \leftarrow !r; \end{aligned}$$

```

    xs ← traverse tl;
    return (x # xs)
done

```

The semantics of a monadic computation t is given by a relation $\llbracket t \rrbracket$, where $(h, h', y) \in \llbracket t \rrbracket$ expresses that if the computation is executed on heap h , then no exception occurs and y and h' are the result value and the new heap, respectively.

By a slight extension of Isabelle’s code generator [13], the monadic terms can be translated to ML (using imperative features) or Haskell (using the ST monad).

1.4 The Catch: Recursive Monadic Definitions

But there is a catch! While much of Isabelle’s reasoning infrastructure can be used for monadic programs as well, the function package cannot cope with functions as simple as *traverse* above.

There are several aspects of the problem. First, the function package looks at the arguments of the function to construct the domain predicate. However, the termination of the function also depends on the heap, which is hidden behind the state monad and not a direct argument of the function.

While this could be solved by breaking the monad abstraction and making the heap a normal argument to the function (which would result in very messy code), the second problem is that *traverse* is inherently partial, as the pointer structure on the heap may be cyclic. Thus, the function package could only produce conditional equations, and one would lose the possibility to use the code generator, which only works for unconditional equations. This limitation of code generation also applies to other partial functions, but in Imperative HOL, where partiality is ubiquitous, it is especially problematic.

The paper on Imperative HOL [7] already observes these shortcomings and provides a workaround using a recursion combinator *MREC* that can express a common case of definitions, including *traverse*. However, the lack of tool support soon becomes a show-stopper as soon as more complex function definitions are needed.

For example, in ongoing work, Bulwahn is formalizing an imperative version of unification. Some of his functions do not fall under the scheme of *MREC*, and had to be defined manually in a tedious and error-prone process. This paper aims to simplify this task by providing a simpler approach and some automation.

1.5 A Solution using Domain Theory

The approach that we take is to abandon well-founded recursion for this task, and resort to domain theory instead, which can express very general recursions over complete partial orders.

The central trick is that the heap monad can be turned into a pointed complete partial order (pcpo) by using the exceptions as a bottom value. Then, any monadic expression built up from pure terms and primitive operations, composed with return and bind operation is continuous by construction. This means that the standard least fixed-point construction can be used to obtain the function, and that the recursive equation can be derived without preconditions.

A well-developed formalization of domain-theoretic concepts is available with Isabelle/HOLCF [19]. While we build on these concepts as a foundation, the constructions are not exposed to the user.

1.6 Related Work

The most closely related work is by Bertot and Komendantsky [5], who use fixed-points in flat ppos to define partial recursive functions in Coq [4], augmented with some classical axioms. They also show that their extension preserves the possibility to extract programs from Coq developments and provide a Coq command that automates the definition process.

This work builds on the same foundations, but goes beyond it in two points:

1. We manage to provide an induction principle, which enables reasoning about the function without having to rely on the underlying iterative construction. This permits induction proofs at a higher level of abstraction, which leads to simpler proofs.
2. By generalizing the approach slightly, we do not only deal with the flat domain of the option type, but can also handle other situations such as the heap monad. This elegantly solves the problem of making recursive definitions of programs in Imperative HOL.

Using the monad abstraction to encapsulate partiality is also not new: In the context of constructive type theory, Capretta [8] and Megacz [18] describe monads that model non-terminating computations coinductively. In a classical logic like Isabelle/HOL, however, the much simpler option monad, which simply adds an extra element to express undefinedness, is sufficient.

The rest of this paper is structured as follows. We first introduce the basic preliminaries from Isabelle/HOLCF (Sect. 2). Then we show how recursive function definitions can be automated in the option monad, which is the simplest setting for our approach (Sect.3). Sect. 4 discusses the automated generation of induction rules from the general fixed-point induction principle. Then we move back to the imperative heap monad (Sect. 5), which is the original motivation for this work. It will be seen that the technique generalizes easily to that more interesting case, and we present a more realistic example. We compare the method to the domain-predicate-based approach and discuss limitations and other issues in Sect. 6.

2 Isabelle/HOL and Isabelle/HOLCF

Isabelle/HOLCF is a definitional extension of Isabelle/HOL with domain-theoretical concepts from the LCF system [11]. Originally developed by Regensburger [22, 21], its design was later improved by Müller et al. [19], notably by the consequent use of type classes. In recent years, the library has been maintained and extended by Huffman [14, 15].

HOLCF defines a type class of complete partial orders (cpo) \sqsubseteq , based on which the standard notions of chain, (least) upper bounds, and continuity are defined as in Fig. 1. Note that the expression $\bigsqcup i. Y i$ only denotes a meaningful value if a least upper bound actually exists. Otherwise its value is arbitrary. A cpo is *pointed* if it has a least element, written \perp . Pointed cpos are also called ppos.

One of the basic results of domain theory is a fixed-point theorem, proving that continuous functions always have a fixed-point that can be reached by iteration:

$$\text{cont } F \implies \text{fixp } F = F (\text{fixp } F) \quad (\text{FIXP})$$

Here, $\text{fixp } F = (\bigsqcup i. F^i \perp)$, and F^i denotes iterated function application.

As one of its main features, HOLCF then introduces a type of continuous functions, written $\alpha \rightarrow \beta$. While this helps to automate many continuity proofs by turning them into type checking,

$$\begin{aligned}
& \text{chain} :: (\text{nat} \Rightarrow \alpha) \Rightarrow \text{bool} \\
& \text{chain } Y \longleftrightarrow (\forall i. Y\ i \sqsubseteq Y\ (\text{Suc } i)) \\
& \text{range} :: (\alpha \Rightarrow \beta) \Rightarrow \beta \text{ set} \\
& \text{range } f = \{y. \exists x. y = f\ x\} \\
& \triangleleft :: \alpha \text{ set} \Rightarrow \alpha \Rightarrow \text{bool} \\
& S \triangleleft x \longleftrightarrow (\forall y. y \in S \longrightarrow y \sqsubseteq x) \\
& \triangleleft\triangleleft :: \alpha \text{ set} \Rightarrow \alpha \Rightarrow \text{bool} \\
& S \triangleleft\triangleleft x \longleftrightarrow S \triangleleft x \wedge (\forall u. S \triangleleft u \longrightarrow x \sqsubseteq u) \\
& \bigsqcup :: \alpha \text{ set} \Rightarrow \alpha \\
& (\bigsqcup i. Y\ i) = (\text{THE } x. \text{range } Y \triangleleft\triangleleft x) \\
& \text{cont} :: (\alpha \Rightarrow \beta) \Rightarrow \text{bool} \\
& \text{cont } f \longleftrightarrow (\forall Y. \text{chain } Y \longrightarrow \text{range } (\lambda i. f\ (Y\ i)) \triangleleft\triangleleft f\ (\bigsqcup i. Y\ i))
\end{aligned}$$

Figure 1: Basic definitions of HOLCF

it also destroys the compatibility with regular Isabelle/HOL developments. Since we are trying to simplify function definitions in HOL, we do not use the continuous function space.

More generally, while this work uses concepts from HOLCF, it is important to note that this is completely transparent and just part of the internal construction. A user of our tool does not have to know domain theory or its formalization in HOLCF.

3 Recursion in the Option Monad

This section shows how to define partial functions in the option monad. It mainly recalls the standard fixed-point construction also used by Bertot and Komendantsky [5], and shows how it is automated in Isabelle/HOL. Later we will generalize it to the heap monad.

We start from the standard option type in Isabelle/HOL, together with the monad operations:

$$\begin{aligned}
& \text{datatype } \alpha \text{ option} = \text{None} \mid \text{Some } \alpha \\
& \text{return } x = \text{Some } x \\
& \text{None} \gg= f = \text{None} \\
& \text{Some } y \gg= f = f\ y
\end{aligned}$$

This monad is known to Haskell programmers as the Maybe monad and models computations with failure. However, it can also be regarded as a (flat) pcpo, where $\perp = \text{None}$.

This basically (ab)uses *None* as the result of a non-terminating computation. The fixed-point law can thus be used to solve recursive equations $f\ x = F\ f\ x$, provided that the functional involved is continuous:

1. Prove that the functional F is continuous.
2. Define $f = \text{fixp } F$.
3. Conclude the equation $f\ x = F\ f\ x$ using the fixed-point theorem (**FIXP**).

$$\begin{array}{ll}
(\bigwedge y. \text{cont } (\lambda x. f x y)) \Longrightarrow \text{cont } f & \text{(LAM)} \\
\text{cont } f \Longrightarrow (\bigwedge y. \text{cont } (g y)) \Longrightarrow \text{cont } (\lambda x. \text{do } y \leftarrow f x; g y x \text{ done}) & \text{(BIND)} \\
\text{cont } (\lambda x. c) & \text{(CONST)} \\
\text{cont } (\lambda f. f x) & \text{(REC)} \\
\text{cont } f \Longrightarrow \text{cont } g \Longrightarrow \text{cont } (\lambda x. \text{if } b \text{ then } f x \text{ else } g x) & \text{(IF)}
\end{array}$$

Figure 2: Continuity rules

Now the primary observation is that if the function is written in monadic style, continuity holds by construction and can be proved automatically following the term structure using the rules given in Fig. 2.

Example 1. As an artificial example, assume some fixed function $\text{step} :: \text{nat} \Rightarrow \text{nat}$ and assume that we want to define a function $\text{trace} :: \text{nat} \Rightarrow \text{nat list}$ that iterates the step function, until it returns zero, keeping all even values in a list. Here is how the function could be written in ML.

```

fun trace n =
  if n = 0 then []
  else if even n then n :: trace (step n)
  else trace (step n)

```

The function is partial and asserting this equation directly in Isabelle/HOL would be unsound. However, we can define its monadic counterpart with return type nat list option : (Note that $\#$ is Isabelle's way of spelling the constructor for non-empty lists)

```

trace n =
  (if n = 0 then return []
   else do tl ← trace (step n);
   (if even n then return (n # tl) else return tl)
  done)

```

The following step-by-step proof shows that continuity of the functional is easily proved in a completely syntax-directed way. The proof obligation is as follows.

1. $\text{cont } (\lambda \text{trace } n.$

$$\begin{array}{l}
\text{if } n = 0 \text{ then return } [] \\
\text{else do } tl \leftarrow \text{trace } (\text{step } n); \\
\text{(if even } n \text{ then return } (n \# tl) \text{ else return } tl) \\
\text{done)}
\end{array}$$

We first move the lambda bound argument out using rule (LAM):

1. $\bigwedge n. \text{cont } (\lambda \text{trace}. \text{if } n = 0 \text{ then return } []$

$$\begin{array}{l}
\text{else do } tl \leftarrow \text{trace } (\text{step } n); \\
\text{(if even } n \text{ then return } (n \# tl) \text{ else return } tl) \\
\text{done)}
\end{array}$$

Applying rule (IF), we obtain two subgoals:

1. $\bigwedge n. \text{cont } (\lambda \text{trace}. \text{return } [])$
2. $\bigwedge n. \text{cont } (\lambda \text{trace}. \text{do } tl \leftarrow \text{trace } (\text{step } n);$

$$\begin{array}{l}
\text{(if even } n \text{ then return } (n \# tl) \text{ else return } tl) \\
\text{done)}
\end{array}$$

The first goal is trivial as it contains no recursive call, and can be discharged with rule (CONST). The other goal contains a bind, and we decompose it using rule (BIND):

1. $\bigwedge n. \text{cont } (\lambda \text{trace}. \text{trace } (\text{step } n))$
2. $\bigwedge n \text{ tl}. \text{cont } (\lambda \text{trace}. \text{if even } n \text{ then return } (n \# \text{tl}) \text{ else return } \text{tl})$

Now, the first goal is a recursive call, and we apply rule (REC). The other goal is again trivially solved using (CONST).

4 Induction Rules for Partial Correctness

Defining the function and deriving the recursive equation is always just half of the problem, since one wants to use induction to reason about the function. For total functions, the induction principle is a consequence of the termination of the function. For partial functions, the function package can generate a similar rule, using the domain predicate as a guard. In this section, we show how to derive an induction rule for partial functions constructed as least fixed-points.

HOLCF provides a general fixed-point induction rule:

$$\text{adm } P \Longrightarrow \text{cont } F \Longrightarrow P \perp \Longrightarrow (\bigwedge f. P f \Longrightarrow P (F f)) \Longrightarrow P (\text{fix } F)$$

Besides continuity, which is already proved at definition time, the property P must hold for \perp and it must be *admissible*, which means that it can be transferred from chains to least upper bounds:

$$\text{adm } P = (\forall Y. \text{chain } Y \longrightarrow (\forall i. P (Y i)) \longrightarrow P (\bigsqcup i. Y i))$$

While this general rule can be used to reason about the function, it is somewhat abstract and not very convenient to use. In particular, the admissibility condition must always be proved when applying the rule. Although HOLCF can automate such proofs in some cases, we would prefer to hide these inconvenient parts of domain theory completely.

4.1 Restriction to partial correctness

It turns out that there is an instance of the general rule which is easier to work with.

If we restrict ourselves to partial correctness properties, i.e., showing that the result of the function, when it is defined, satisfies some predicate, then matters become straightforward. More precisely, we replace the predicate P with the instance $\lambda f. \forall x y. f x = \text{Some } y \longrightarrow Q x y$. Then the admissibility condition can be discharged once and for all, since this instance is always admissible.

Thus, if f is the recursive function, and F is the corresponding functional, the following rule can be derived.

$$\frac{\bigwedge f x y. (\bigwedge z r. f z = \text{Some } r \Longrightarrow Q z r) \Longrightarrow F f x = \text{Some } y \Longrightarrow Q x y}{f x = \text{Some } y \Longrightarrow Q x y}$$

Note that the statement of this rule makes no mentioning of the iterative fixed-point construction. Presenting this rule to the user hides the details of this construction, which allows reasoning on a more abstract level.

Example 2. For example, the instance for *trace* is as follows:

$$\begin{array}{l}
 \bigwedge \text{trace } n \text{ } ys. \\
 (\bigwedge z r. \text{trace } z = \text{Some } r \implies Q \ z \ r) \implies \\
 (\text{if } n = 0 \text{ then return } [] \\
 \text{else do } tl \leftarrow \text{trace } (\text{step } n); \\
 \text{if even } n \text{ then return } (n \# tl) \text{ else return } \\
 tl) \\
 \text{done}) = \\
 \text{Some } ys \implies \\
 Q \ n \ ys \\
 \hline
 \text{trace } n = \text{Some } ys \implies Q \ n \ ys
 \end{array}$$

4.2 Induction Rule Refinement

The raw induction rule as shown above can still be improved. First, the control flow in the definition gives rise to three cases, one for $n = 0$, one for *even* n , and one for \neg *even* n . Second, the sequencing using $\gg=$ can be decomposed, since the whole expression is only defined when all relevant subexpressions are defined. Moreover, the induction hypothesis is likely to be only useful to prove that the recursive call satisfies the property Q . The rule that a user would like to see is roughly the following:

$$\begin{array}{l}
 Q \ 0 \ [] \\
 \bigwedge n \ tl. n \neq 0 \implies Q \ (\text{step } n) \ tl \implies \text{even } n \implies Q \ n \ (n \\
 \# \ tl) \\
 \bigwedge n \ tl. n \neq 0 \implies Q \ (\text{step } n) \ tl \implies \neg \text{even } n \implies Q \ n \\
 tl \\
 \hline
 \text{trace } n = \text{Some } ys \implies Q \ n \ ys
 \end{array}$$

To arrive at this simpler form, the following steps are necessary:

1. Decompose the program structure by splitting the function body into smaller steps:
 - A premise $(t \gg= (\lambda x. f \ x)) = \text{Some } y$ is replaced by $t = \text{Some } x$ and $f \ x = \text{Some } y$
 - A premise $\text{Some } t = \text{Some } y$ (which arises from a return statement) is replaced by $t = y$.
 - Conditionals like $(\text{if } b \text{ then } x \text{ else } x') = \text{Some } y$ are split up into two cases, with premises b and $\neg b$.
2. Use the induction hypothesis to replace premises of the form $f \ z = \text{Some } r$ by $Q \ z \ r$. When all occurrences of f are replaced, the general induction hypothesis can be discarded.
3. Clean up the context by substituting premises of the form $v = t$ where v is a variable.

5 Recursion in the Heap Monad

We now move from the option monad to the more interesting heap monad. In fact, not much of the process has to be adapted.

The heap pcpo. Unlike the option pcpo, the heap pcpo is not flat, since its values represent state transformations. The order is defined as $Heap\ f \sqsubseteq Heap\ g \iff (\forall h. f\ h = bot \vee f\ h = g\ h)$, where $bot = (Inr\ Exn, h_0)$ for some arbitrary but fixed heap h_0 . This implies $\perp = Heap\ (\lambda h. bot)$.

Recursive definitions. After proving continuity of $\gg=$, which is tedious but straightforward, the definition process remains the same as for the option monad. We automatically prove continuity of the functional by applying the rules from Fig. 2 in a syntax-directed way. After that, the function can be defined as a fixed-point.

Induction rule generation. In the induction rule, the condition $f\ x = Some\ y$ is replaced with its counterpart for heap-manipulating programs, the condition $(h, h', y) \in \llbracket f\ x \rrbracket$ (cf. Sect. 1.3). The inductive property Q now also refers to the heap before and after the computation. As in the option case, we must prove that this partial correctness property is always admissible:

$$adm\ (\lambda f. \forall x\ h\ h'\ y. (h, h', y) \in \llbracket f\ x \rrbracket \implies Q\ x\ h\ h'\ y)$$

Induction rule refinement. In the refinement process, the program structure is decomposed as described above. For the primitive heap operations, additional refinement steps are added, which replace them by their counterparts with explicit heap. For example, the premise $(h, h', y) \in \llbracket !r \rrbracket$ is replaced by $y = get-ref\ r\ h$ and $h' = h$.

Example 3. The refined induction rule for *traverse* has the following form:

$$\frac{\bigwedge h'. Q\ Empty\ h'\ h'\ [] \quad \bigwedge h_1\ h_2\ x'\ r\ n. Q\ (get-ref\ r\ h_1)\ h_1\ h_2\ n \implies Q\ (Node\ x'\ r)\ h_1\ h_2\ (x' \# n)}{(h, h', y) \in \llbracket traverse\ x \rrbracket \implies Q\ x\ h\ h'\ y}$$

Example 4. We now discuss a more realistic example that arises in the formalization of an imperative unification algorithm mentioned previously. We refer to Baader and Nipkow [1, ch. 4.8] for a textbook description of imperative unification. To avoid expensive allocations, the algorithm keeps track of substitutions by directly updating references in the terms themselves.

In the formalization, heap-allocated mutable terms consist of variables, constants and binary applications:

datatype $\alpha\ rtrm = Var\ \alpha\ (\alpha\ rtrm\ ref\ option) \mid Const\ \alpha \mid App\ (\alpha\ rtrm\ ref)\ (\alpha\ rtrm\ ref)$

The type argument α is only used for names. Note that variables carry an optional reference cell, which is used to mark that a variable has already been assigned some other term. Applying a substitution for that variable only requires an update of the relevant reference, which also affects other occurrences of the same variable, since the reference is shared. A value of *None* means that the variable is unassigned.

We will only show a simple part of the unification algorithm, namely the function *occurs*, which checks if a variable r_1 appears in some term r_2 :

```

occurs r1 r2 =
do t ← !r2;
(case t of
  Var n σ ⇒
    if r1 = r2 then return True
    else case σ of None ⇒ return False | Some r' ⇒ occurs r1 r'
  | Const n ⇒ return False
  | App r3 r4 ⇒ do b ← occurs r1 r3;
    (if b then return True else occurs r1 r4)
done)
done

```

This is a simple recursive traversal of r_2 , except that in the variable case, the traversal continues if the variable has already been instantiated. Since the term on the heap may contain cycles the function can diverge. Also note that the check $r_1 = r_2$ is a pointer equality test, not structural equality.

To formulate the correctness property of such a function, it is convenient to re-state the property “ r_1 occurs in r_2 ” as an inductive relation $occurs\text{-}in :: heap \Rightarrow \alpha \text{ rtrm} \Rightarrow \alpha \text{ rtrm} \Rightarrow bool$. Then, the crucial property connects $occurs$ and $occurs\text{-}in$:

$$(h, h', b) \in \llbracket occurs\ r_1\ r_2 \rrbracket \Longrightarrow get\text{-}ref\ r_1\ h = Var\ c\ None \Longrightarrow occurs\text{-}in\ h\ r_1\ r_2 = b$$

Note that there is no assumption that the pointer structures are acyclic, which is implicit in the assumption that the call to $occurs$ terminates and returns b . As there is no structural induction principle that can be used here, and we must use induction over the computation of the function—in other words, fixed-point induction.

The induction rule produced by our prototype implementation is given below. With this rule, the inductive proof of the correctness property is straightforward. Note how the premises of the rule correspond to the cases in the function definition.

$$\begin{array}{l}
\bigwedge r_1\ h\ n\ \sigma. get\text{-}ref\ r_1\ h = Var\ n\ \sigma \Longrightarrow P\ r_1\ r_1\ h\ h\ True \\
\bigwedge r_1\ r_2\ h\ n. r_1 \neq r_2 \Longrightarrow get\text{-}ref\ r_2\ h = Var\ n\ None \Longrightarrow P\ r_1\ r_2\ h\ h\ False \\
\bigwedge r_1\ r_2\ h' y\ h\ n\ r'. \\
\quad r_1 \neq r_2 \Longrightarrow P\ r_1\ r'\ h\ h'\ y \Longrightarrow get\text{-}ref\ r_2\ h = Var\ n\ (Some\ r') \Longrightarrow P\ r_1\ r_2\ h\ h'\ y \\
\bigwedge r_1\ r_2\ h\ n. get\text{-}ref\ r_2\ h = Const\ n \Longrightarrow P\ r_1\ r_2\ h\ h\ False \\
\bigwedge r_1\ r_2\ h' h\ r_3\ r_4\ b. get\text{-}ref\ r_2\ h = App\ r_3\ r_4 \Longrightarrow P\ r_1\ r_3\ h\ h'\ b \Longrightarrow b \Longrightarrow P\ r_1 \\
r_2\ h\ h'\ True \\
\bigwedge r_1\ r_2\ h'' y\ h\ r_3\ r_4\ h'\ b. \\
\quad get\text{-}ref\ r_2\ h = App\ r_3\ r_4 \Longrightarrow \\
\quad P\ r_1\ r_3\ h\ h'\ b \Longrightarrow \neg b \Longrightarrow P\ r_1\ r_4\ h'\ h''\ y \Longrightarrow P\ r_1\ r_2\ h\ h''\ y \\
\hline
(y, r_1, r_2) \in \llbracket occurs\ h\ h' \rrbracket \Longrightarrow P\ h\ h'\ y\ r_1\ r_2
\end{array}$$

While this rule looks intimidating at first, our (limited) practical experience suggests that there is no way around it. Before automation was available, the definition of the function and the proof of a similar induction rule took about 450 lines of very technical proof script, which is more than the correctness proof itself. The situation for the rest of the imperative unification algorithm is similar. With our new approach, these manual proofs are no longer needed.

6 Discussion

Implementation. The implementation of our technique is still in a prototype stage. It works well with the examples like the ones presented in this paper but lacks several user-friendly features that would be needed for productive use, e.g., support for mutually recursive definitions and pattern matching. Moreover, the HOLCF dependencies should be reduced to a minimum.

Higher-order recursion. Currently, our approach expects a fixed set of constructs in monadic terms: $\gg=$, conditionals, recursive calls, and constant expressions not involving recursive calls (cf. the continuity rules in Fig. 2). This basically limits the functions that can be defined to a first-order fragment, since recursive calls must be applied and cannot be passed to functionals like *map* or *fold*. To support higher-order recursion, one must handle more constructs, such as a monadic map combinator *mapM*. Then, a suitable continuity rule must be known to the tool, and (optionally) a rule to be used in the induction rule refinement step. The format of such rules is not yet entirely clear, and it is part of future work to see if this extension is possible and helpful in practice.

Option type vs. domain predicate. With the ability to define non-terminating recursive functions in the option monad, we effectively have a new and alternative technique for defining partial functions. This raises the question about the relative merits of the two approaches, and whether one should be preferred over the other. While this question can only be answered by comparing the approaches in concrete applications, a few general things can be said:

- The main difference between the techniques is the format of the function itself: while the function package produces a total function and a predicate *dom* that describes the arguments x where $f x$ “makes sense”, the option method produces a function that returns an option type, where partiality is explicit in the result. As a price for this more accurate model, the recursive equation must be written monadically, to deal with the bottom values that may arise in recursive calls.
- The function package constrains the recursive equations with the domain predicate, whereas the option method produces unconditional equations. Thus, partial functions defined using the option method can be used with Isabelle’s code generator.
- The function package provides special support for tail-recursive definitions. If the function is tail-recursive, the unconditional specification can be derived as a theorem even for partial functions.

The same functionality could be provided using the domain-theoretic approach: We choose the identity monad (i.e., no monad at all), and take the flat pcpo where $x \sqsubseteq y = (x = u \vee x = y)$ where u is an arbitrary but fixed value. Since the bind operation (which is just application) is not continuous, we cannot define arbitrary functions. We can however define tail-recursive functions, since they can be written without bind. Thus, tail-recursion can be seen as a special case of the monadic approach.

Other Applications. Up to now, only option and state-exception monads are supported by our approach. The question arises whether our treatment can be transferred to recursive definitions in other monads, e.g., continuation or resumption monads. This is the subject of future work.

However, even in the current state, the approach can be of use in scenarios other than Imperative HOL: For example, Thiemann and Sternagel [24] use an error monad to formalize a simple XML parser. Termination of the parser is irrelevant to the rest of their development, so it is currently assumed as an axiom, in order to obtain unconditional equations. Using our approach, these equations should come for free, but without axioms.

7 Conclusion

Monadic functions present a challenge to the automated definition mechanisms based on well-founded recursion. We have shown that by using the fixed-point theorem for complete partial orders, such definitions can be made with surprising ease and result in equations that need no domain condition. Automatically generated custom induction rules make the resulting functions convenient to use, without having to refer to the iterative construction used internally.

While it is perhaps not surprising that domain theory is up to this task, using it for monadic definitions is particularly convenient, since conditions like continuity and admissibility hold by construction, and the overhead of handling undefinedness is absorbed by the monad.

Acknowledgments I want to thank Lukas Bulwahn for many interesting discussions about improving tool support for Imperative HOL, and for providing helpful feedback on a draft of this paper and the implementation.

References

- [1] F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.
- [2] G. Barthe, J. Forest, D. Pichardie, and V. Rusu. Defining and reasoning about recursive functions: a practical tool for the Coq proof assistant. In M. Hagiya and P. Wadler, editors, *Functional and Logic Programming (FLOPS 2006)*, volume 3945 of *Lecture Notes in Computer Science*, pages 114 – 129. Springer Verlag, 2006.
- [3] S. Berghofer, T. Nipkow, C. Urban, and M. Wenzel, editors. *Theorem Proving in Higher Order Logics (TPHOLs 2009)*, 22nd International Conference, Munich, Germany, August 2009, *Proceedings*, volume 5674 of *Lecture Notes in Computer Science*. Springer Verlag, 2009.
- [4] Y. Bertot and P. Castéran. *Interactive theorem proving and program development: Coq’Art: the calculus of inductive constructions*. Texts in theoretical computer science. Springer Verlag, 2004.
- [5] Y. Bertot and V. Komendantsky. Fixed point semantics and partial recursion in Coq. In *Principles and practice of declarative programming (PPDP ’08)*, pages 89–96, New York, NY, USA, 2008. ACM.
- [6] A. Bove and V. Capretta. Modelling general recursion in type theory. *Mathematical Structures in Computer Science*, 15(4):671–708, 2005.
- [7] L. Bulwahn, A. Krauss, F. Haftmann, L. Erkök, and J. Matthews. Imperative functional programming in Isabelle/HOL. In O. Ait Mohamed, C. Muñoz, and S. Tahar, editors, *Theorem Proving in Higher Order Logics (TPHOLs 2008)*, volume 5170 of *Lecture Notes in Computer Science*, pages 134–149. Springer Verlag, 2008.
- [8] V. Capretta. General recursion via coinductive types. *Logical Methods in Computer Science*, 1(2):1–18, 2005.
- [9] R. L. Constable and N. P. Mendler. Recursive definitions in type theory. In R. Parikh, editor, *Logic of Programs*, volume 193 of *Lecture Notes in Computer Science*, pages 61–78. Springer Verlag, 1985.

- [10] C. Dubois and V. Donzeau-Gouge. A step towards the mechanization of partial functions: domains as inductive predicates. In *CADE-15 Workshop on mechanization of partial functions*, 1998.
- [11] M. J. C. Gordon, R. Milner, and C. P. Wadsworth. *Edinburgh LCF: A Mechanised Logic of Computation*, volume 78 of *Lecture Notes in Computer Science*. Springer Verlag, 1979.
- [12] D. Greve. Assuming termination. In *ACL2 Workshop Proceedings*, 2009.
- [13] F. Haftmann and T. Nipkow. Code generation via higher-order rewrite systems. In *Functional and Logic Programming (FLOPS 2010)*, LNCS. Springer, 2010. To appear.
- [14] B. Huffman. Reasoning with powerdomains in Isabelle/HOLCF. In O. Ait Mohamed, C. Muñoz, and S. Tahar, editors, *TPHOLs 2008: Emerging Trends Proceedings*, pages 45–56, 2008. Department of Electrical and Computer Engineering, Concordia University.
- [15] B. Huffman. A purely definitional universal domain. In Berghofer et al. [3], pages 260–275.
- [16] A. Krauss. *Automating Recursive Definitions and Termination Proofs in Higher-Order Logic*. PhD thesis, Institut für Informatik, Technische Universität München, Germany, 2009.
- [17] A. Krauss. Partial and nested recursive function definitions in higher-order logic. *Journal of Automated Reasoning*, 44(4):303–336, 2010.
- [18] A. Megacz. A coinductive monad for prop-bounded recursion. In A. Stump and H. Xi, editors, *Proceedings of the ACM Workshop Programming Languages meets Program Verification, PLPV 2007*, pages 11–20, New York, NY, USA, 2007. ACM.
- [19] O. Müller, T. Nipkow, D. von Oheimb, and O. Slotosch. HOLCF=HOL+LCF. *Journal of Functional Programming*, 9(2):191–223, 1999.
- [20] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science*. Springer Verlag, 2002.
- [21] F. Regensburger. *HOLCF: Eine konservative Erweiterung von HOL um LCF*. PhD thesis, Technische Universität München, 1994.
- [22] F. Regensburger. HOLCF: higher order logic of computable functions. In E. T. Schubert, P. J. Windley, and J. Alves-Foss, editors, *TPHOLs*, volume 971 of *Lecture Notes in Computer Science*, pages 293–307. Springer Verlag, 1995.
- [23] K. Slind. Function definition in Higher-Order Logic. In J. von Wright, J. Grundy, and J. Harrison, editors, *Theorem Proving in Higher Order Logics (TPHOLs '96)*, volume 1125 of *Lecture Notes in Computer Science*, pages 381–397. Springer Verlag, 1996.
- [24] R. Thiemann and C. Sternagel. Certification of termination proofs using CeTA. In Berghofer et al. [3], pages 452–468.