Embedding the refinement calculus in Coq

João Alpuim [a], Wouter Swierstra [b]

[a] *University of Hong Kong, Hong Kong*
[b] *Universiteit Utrecht, Netherlands*

A B S T R A C T

The *refinement calculus* and *type theory* are both frameworks that support the specification and verification of programs. This paper presents an embedding of the refinement calculus in the interactive theorem prover Coq, clarifying the relation between the two. As a result, refinement calculations can be performed in Coq, enabling the interactive calculation of formally verified programs from their specification.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

The idea of *deriving* a program from its specification can be traced back to [10,12,16]. The *refinement calculus* [5,17,3] defines a formal methodology that can be used to construct a derivation of a program from its specification step by step. Crucially, the refinement calculus presents single language for describing both programs and specifications.

Deriving complex programs using the refinement calculus is no easy task. The proofs and obligations can quickly become too complex to manage by hand. Once you have completed a derivation, the derived program must still be transcribed to a programming language in order to execute it – a process which can be rather error-prone [17, Chapter 19].

To address both these issues, we show how the refinement calculus can be embedded in Coq, an interactive proof assistant based on dependent types. Although others have proposed similar formalizations of the refinement calculus [4,13], this paper presents the following novel contributions:

- After giving a brief overview of the refinement calculus (Section 2), we begin by developing a library of predicate transformers in Coq, based on indexed containers [1,13], making extensive use of dependent types (Section 3). We define a *refinement relation*, corresponding to a morphism between indexed containers, enabling us to prove several simple refinement laws in Coq.
- Next we show how to embed effects, such as mutable state and general recursion, in Coq using a free monad. We assign semantics to programs in this monad using the predicate transformers and refinement relation we described previously (Section 4).
- These definitions give us the basic building blocks for formalizing derivations in the refinement calculus. They do, however, require that the derived program is known *a priori*. We address this and other usability issues (Section 5).

- Finally, we validate our results by performing a small case study.[1] In particular, we show how we can use this library to calculate interactively the key ingredients of a data structure for *persistent arrays* (Section 6).

2. Refinement calculus

The refinement calculus, as presented by Morgan [17], extends Dijkstra's Guarded Command language with a new language construct for specifications. The specification [pre, post] is satisfied by a program that, when supplied an initial state satisfying the precondition pre, can be executed to produce a final state satisfying the postcondition post. Crucially, this language construct may be mixed freely with (executable) code constructs.

Besides these specifications, the refinement calculus defines a *refinement relation* between programs, denoted by $p_1 \sqsubseteq p_2$. This relation holds when **forall** P, $\mathsf{wp}(p_1, P) \Rightarrow \mathsf{wp}(p_2, P)$, where wp denotes the usual weakest precondition semantics of a program and its desired postcondition. Intuitively, you may want to read $p_1 \sqsubseteq p_2$ as stating that p_2 is a 'more precise specification' than p_1.

A program is said to be *executable* when it is free of specifications and only consists of executable statements. Morgan [17] refers to such executable programs as *code*. To calculate an executable program C from its specification S, you must find a series of refinement steps, $S \sqsubseteq M_0 \sqsubseteq M_1 \sqsubseteq ... \sqsubseteq C$. Typically, the intermediate programs, such as M_0 and M_1, mix executable code fragments and specifications.

To find such derivations, Morgan [17] presents a catalog of lemmas that can be used to refine a specification to an executable program. Some of these lemmas define when it is possible to refine a specification to code constructs. These lemmas effectively describe the semantics of such constructs. For example, the following law may be associated with the **skip** command:

Lemma 1 (**skip**). *If* pre $\Rightarrow$ post, *then* [pre, post] $\sqsubseteq$ **skip**.

Besides such primitive laws, there are many recurring patterns that pop up during refinement calculations. For example, combining the rules for sequential composition and assignment, the *following assignment* lemma holds:

Lemma 2 (**following assignment**). *For any term E,*

$$[\mathsf{pre}, \mathsf{post}] \sqsubseteq [\mathsf{pre}, \mathsf{post}\,[w \backslash E]]; w ::= E$$

We illustrate how these rules may be used to calculate the definition of a program from its specification. Suppose we would like to swap the values of two variables, x and y. We may begin by formulating the specification of our problem as:

$$[\,x = X \wedge y = Y, x = Y \wedge y = X\,]$$

Using the two lemmas we saw above, we can refine this specification to an executable program. The corresponding calculation is given in Fig. 1. Note that we have chosen to give a simple derivation that contains some redundancy, such as the final **skip** statement, but uses a modest number of auxiliary lemmas and definitions.

For such small programs, these derivations are manageable by hand. For larger or more complex derivations, it can be useful to employ a computer to verify the correctness of the derivation and even assist in its construction. In the coming sections we will develop a Coq library for precisely that.

3. Predicate transformers

In this section, we will assume there is some type S, representing the state that our programs manipulate. In Section 4 we will show how this can be instantiated with a (model of a) heap. For now, however, the definitions of specifications, refinement, and predicate transformers will be made independently of the choice of state.

We begin by defining a few basic constructions in Coq:

> **Definition** Pred $(A : \mathsf{Type}) : \mathsf{Type} := A \to \mathsf{Type}.$

This defines the type Pred A of predicates over some type A. Using this definition we can define a subset relation between predicates as follows:

> **Definition** subset $(A : \mathsf{Type})\,(P_1\ P_2 : \mathsf{Pred}\,A) := \mathbf{forall}\ x, P_1\ x \to P_2\ x.$

---

[1] All the code and examples presented in this paper can be found online at https://github.com/jalpuim/dtp-refinement.

$$[\,x = X \land y = Y, x = Y \land y = X\,]$$
$$\sqsubseteq \quad \{\text{ by the following assignment law }\}$$
$$[\,x = X \land y = Y, t = Y \land y = X\,]; x ::= t$$
$$\sqsubseteq \quad \{\text{ by the following assignment law }\}$$
$$[\,x = X \land y = Y, t = Y \land x = X\,]; y ::= x; x ::= t$$
$$\sqsubseteq \quad \{\text{ by the following assignment law }\}$$
$$[\,x = X \land y = Y, y = Y \land x = X\,]; t ::= y; y ::= x; x ::= t$$
$$\sqsubseteq \quad \{\text{ by the law for skip }\}$$
$$\textbf{skip}; t ::= y; y ::= x; x ::= t$$

Fig. 1. Derivation of the swap program.

A predicate P_1 is a subset of the predicate P_2, if any state satisfying P_1 also satisfies P_2. In the remainder of this paper, we will write $P_1 \subseteq P_2$ when the property subset $P_1\ P_2$ holds.

Next we can define the PT data type, consisting of a precondition and postcondition:

> **Record** PT $(A : \text{Type}) : \text{Type} :=$
> MkPT $\{$ pre : Pred S;
>   post : **forall** $s : S$, pre $s \to A \to$ Pred $S\,\}$.

The postcondition is a relation between the input state, a proof that this input state satisfies the precondition, the value of type A returned by the computation, and the output state. To avoid the need for 'ghost variables', we allow this relation to refer to both the input and output states [19–21]. As this data type will be used to represent specifications, we will use the notation $[\,P, Q\,]$ rather than the more verbose MkPT $P\ Q$.

As its name suggests, the PT type has an obvious interpretation as a *predicate transformer*, i.e., a function mapping predicates to predicates:

> **Definition** semantics $\{A : \text{Type}\}$ $(pt : \text{PT } A) : \text{Pred } (A * S) \to \text{Pred } S :=$
> **fun** $P\ s \Rightarrow \{\,p : \text{pre } pt\ s\ \&\ \textbf{forall } s'\ v, \text{post } pt\ s\ p\ v\ s' \to P\ (v, s')\,\}$.

The semantics function computes the weakest precondition necessary to guarantee that the desired postcondition P holds after executing a program satisfying the given specification pt. Intuitively, the precondition of the specification must hold and the postcondition must imply P. We will sometimes write $[\![pt]\!]$ rather than semantics pt for the sake of brevity. In what follows, we will sometimes leave out implicit argument, such as the argument $\{A : \text{Type}\}$ in the semantics function, from the typeset code presented here.

Next, we characterize the refinement relation between two values of type PT as follows:

> **Inductive** Refines $(pt_1\ pt_2 : \text{PT } A) : \text{Type} :=$
> Refinement : **forall** $(d : \text{pre } pt_1 \subseteq \text{pre } pt_2)$,
> $(\textbf{forall } s\ p\ v, \text{post } pt_2\ s\ (d\ s\ p)\ v \subseteq \text{post } pt_1\ s\ p\ v) \to$
>   Refines $pt_1\ pt_2$.

We consider pt_2 to be a refinement of pt_1 when the precondition of pt_1 implies the precondition of pt_2 and the postcondition of pt_2 implies the postcondition of pt_1. As our postconditions are relations, we need to do some work to describe the latter condition. In particular, we need to transform the assumption that the initial state holds for the precondition of pt_1 to produce a proof that the precondition of pt_2 also holds for the same initial state. To do so, we use the first condition, d, that the precondition of pt_1 implies the precondition of pt_2. We will use the notation, $pt_1 \sqsubseteq pt_2$, for the proposition Refines $pt_1\ pt_2$.

To validate the correctness of this definition, we will show that it satisfies the characterization of refinement in terms of weakest precondition semantics given in Section 2. To do so, we have proven the following soundness result:

> **Theorem** soundness : **forall** $(pt_1\ pt_2 : \text{PT } A)$,
> $pt_1 \sqsubseteq pt_2 \leftrightarrow \textbf{forall } P, [\![pt_1]\!]\ P \subseteq [\![pt_2]\!]\ P$.

In other words, the Refines relation adheres to the characterization of the refinement relation in terms of predicate transformer semantics. The proof is almost trivial after unfolding the various definitions involved.

Even if we have not yet fixed the state space S, we can already prove that the structural laws of the refinement calculus, such as strengthening of postconditions, hold:

Lemma strengthen $(P : \text{Pred } S) (Q_1 \ Q_2 : \textbf{forall } s, P \ s \rightarrow A \rightarrow \text{Pred } S)$
$\quad (H : \textbf{forall } (s : S) (v : A) (p : P \ s), Q_1 \ s \ p \ v \subseteq Q_2 \ s \ p \ v) :$
$\quad \text{Refines } ([P, Q_2 \ ]) ([P, Q_1 \ ]).$

To prove this lemma, we need to show that $P \subseteq P$ and that the postcondition Q_1 implies Q_2. The first proof is trivial; the second follows immediately from our hypothesis. Similarly, we can show that the refinement relation is both transitive and reflexive.

These definitions by themselves are not very useful. Before we can perform any *program* derivation, we first need to fix our *programming language*.

4. Embedding in Coq

In our previous work [22], we showed how to define a *deep embedding* of small, imperative programming language. By defining predicate transformers associated with the various syntactic constructs, we could then calculate a program from its specification. Such a deep embedding of the language allows us to inspect the abstract syntax of our object language, enabling us to generate programs calculated from their specifications. This approach does have its drawbacks. In particular, the imperative programming language was too restrictive to cover many, more interesting examples. Furthermore, the mutable references were only allowed to store integers. In this extended paper, we address these issues.

To begin with, we fix our choice of state S to be a finite map from an abstract type Ptr to some type of values v. All our developments are parametrized over the particular choice of type to store on the heap. Although we use the finite map modules from Coq's standard library to model the heap, several alternative representations exist [18,23,21].

Instead of a deep embedding of the imperative language, we will provide a deep embedding only of the *effects* of our object language, namely mutable state and unbounded recursion by defining a suitable *free monad*. We will then piggyback on Coq's programming language, Gallina, to define programs using these effects. We can then customize Coq's extraction to OCaml to map these syntactic constructs to their native OCaml counterparts.

We model the effects as an inductive data type in Coq:

Inductive WhileL $(a : \text{Type}) : \text{Type} :=$
$\quad | \ \text{New} \ : v \rightarrow (\text{Ptr} \rightarrow \text{WhileL } a) \rightarrow \text{WhileL } a$
$\quad | \ \text{Read} : \text{Ptr} \rightarrow (v \rightarrow \text{WhileL } a) \rightarrow \text{WhileL } a$
$\quad | \ \text{Write} : \text{Ptr} \rightarrow v \rightarrow \text{WhileL } a \rightarrow \text{WhileL } a$
$\quad | \ \text{While} : (S \rightarrow \text{Prop}) \rightarrow (S \rightarrow \text{bool}) \rightarrow (\text{WhileL unit}) \rightarrow$
$\quad\quad\quad\quad \text{WhileL } a \rightarrow \text{WhileL } a$
$\quad | \ \text{Spec} \ : \text{PT } a \rightarrow \text{WhileL } a$
$\quad | \ \text{Return} : a \rightarrow \text{WhileL } a.$

This data type has constructors for the creation of (New), access of (Read), and assignment to (Write) mutable references. Each of these constructors takes a continuation as argument, representing the remaining computation to perform. Loops may be introduced using the While constructor that takes four arguments: the loop invariant of type $S \rightarrow \text{Prop}$; the condition of type $S \rightarrow \text{bool}$; the loop body of type WhileL unit; and the remaining computation. The constructor Spec contains the specification of an unfinished program fragment. The refinement laws we will define shortly determine how such specifications may be refined to executable code. Finally, the Return constructor simply returns a 'pure' Coq value.

4.1. Semantics

Before discussing the refinement calculation further, we need to fix the semantics of our language. We shall do so by associating a predicate transformer, i.e., a value of type PT, with every constructor of the WhileL data type.
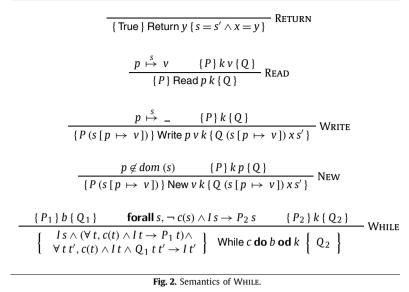
Each rule in Fig. 2 associates pre- and postconditions, i.e., a value of type PT, with the constructors of the WhileL data type. We use the somewhat suggestive notation, $\{P\} \ c \ \{Q\}$ to associate with the statement c the conditions $[P, Q]$. We also use the notation $p \overset{s}{\mapsto} v$ to denote that s maps the pointer p to value v, and $p \overset{s}{\mapsto} \_$ to denote that p has some associated value in s. Finally, the notation $P(s[p \mapsto v])$ denotes that the condition P should hold after updating the state s, mapping the pointer p to value v.

These rules are not added as axioms to Coq; nor are they the constructors of an inductive data type. Rather, we can assign semantics to our WhileL data type directly, as a recursive function:

Fixpoint semantics $(c : \text{WhileL } a) : \text{PT } a$

In addition to the rules from Fig. 2, this function simply maps specifications, represented by the Spec constructor, to their associated predicate transformer.

Let us examine the rules in Fig. 2 a bit more closely. Each precondition may refer to an initial state s; each postcondition is formulated as a relation between an initial state s, satisfying the precondition, the result returned x, and the final state s'.

$$\frac{}{\{\,\text{True}\,\}\,\text{Return}\,y\,\{\,s = s' \wedge x = y\,\}}\ \text{RETURN}$$

$$\frac{p \stackrel{s}{\mapsto} v \qquad \{P\}\,k\,v\,\{Q\}}{\{P\}\,\text{Read}\,p\,k\,\{Q\}}\ \text{READ}$$

$$\frac{p \stackrel{s}{\mapsto} \_ \qquad \{P\}\,k\,\{Q\}}{\{P\,(s\,[\,p \mapsto v\,])\,\}\,\text{Write}\,p\,v\,k\,\{Q\,(s\,[\,p \mapsto v\,])\,x\,s'\,\}}\ \text{WRITE}$$

$$\frac{p \notin dom\,(s) \qquad \{P\}\,k\,p\,\{Q\}}{\{P\,(s\,[\,p \mapsto v\,])\,\}\,\text{New}\,v\,k\,\{Q\,(s\,[\,p \mapsto v\,])\,x\,s'\,\}}\ \text{NEW}$$

$$\frac{\{P_1\}\,b\,\{Q_1\} \qquad \textbf{forall}\ s, \neg\,c(s) \wedge I\,s \to P_2\,s \qquad \{P_2\}\,k\,\{Q_2\}}{\left\{\begin{array}{l} I\,s \wedge (\forall\,t, c(t) \wedge I\,t \to P_1\,t) \wedge \\ \forall\,t\,t', c(t) \wedge I\,t \wedge Q_1\,t\,t' \to I\,t' \end{array}\right\}\ \text{While}\,c\,\textbf{do}\,b\,\textbf{od}\,k\,\left\{\ Q_2\ \right\}}\ \text{WHILE}$$

Fig. 2. Semantics of WHILE.

For example, the postcondition of the Return rule states that the initial state s is equal to the final state s' and the result of the computation is indeed the argument of the Return constructor.

The other rules are defined by induction on the program. For example, the rule for Read states that, provided the current state maps the pointer p to the value v, and the remainder of the program has an associated precondition P and postcondition Q, the composite program has the same pre- and postcondition. When the state changes, such as in the rules for Write and New, this style of definition can be a bit confusing. The Write rule, for instance, computes the precondition and postcondition associated with the remaining program k. The composite program Write $p\,v\,k$ then requires that the precondition P associated with k holds *after* updating the state. Similarly, the postcondition Q should relate the state after being updated with the result x and final state s'. The New rule follows exactly the same pattern.

Finally, the WHILE rule is the most complex. Besides the precondition, P_1, and postcondition, Q_1, associated with the body of the loop b, the WHILE rule requires the programmer to specify the loop invariant, I, and continuation k. The precondition of the WHILE rule consists of three conjuncts:

- the invariant I must hold initially;
- the boolean guard c and the invariant must together imply the precondition of the loop body;
- the loop body must preserve the invariant.

Furthermore, when the loops is finished and c no longer holds, this must imply the precondition of the continuation k. Finally, the entire expression has the postcondition Q_2 associated with the continuation k. Note that this formulation captures *partial correctness*; there is no variant ensuring that the loop must terminate eventually.

Using these semantics, we now define a refinement relation between statements programs in the WhileL language:

> **Definition** WhileRefines $(c_1\,c_2 : \text{WhileL}\,a)$
> := Refines (semantics c_1) (semantics c_2).

Once again, we will use the notation $c_1 \sqsubseteq c_2$ when WhileRefines $c_1\,c_2$ holds.

4.2. Composing programs

It is rather straightforward to show that WhileL supports a monadic bind operation:

> **Fixpoint** bind $\{a\,b : \text{Type}\}\,(w : \text{WhileL}\,a)\,(k : a \to \text{WhileL}\,b) : \text{WhileL}\,b$
> := **match** w **with**
> | New $v\,c \Rightarrow$ New v (**fun** $p \Rightarrow$ bind $(c\,p)\,k$)
> | Read $p\,c \Rightarrow$ Read p (**fun** $v \Rightarrow$ bind $(c\,v)\,k$)
> | Write $p\,v\,c \Rightarrow$ Write $p\,v$ (bind $c\,k$)
> | While *Inv cond body* $c \Rightarrow$ While *Inv cond body* (bind $c\,k$)
> | Spec $pt \Rightarrow$ Spec (*bindPT pt* (**fun** $x \Rightarrow$ semantics $(k\,x)$))
> | Return $x \Rightarrow k\,x$
> **end**.

The function is inductively defined on the first program w. The first four cases simply propagate the bind operation to their respective continuations. The only interesting case is that for specifications. The Spec case creates a new PT, by using its associated predicate transformer composition operator *bindPT*, corresponding to taking the relational composition of the specification and the semantics associated with the continuation. Finally, the Return case applies k to the value x.

Using this definition of bind, we can assemble larger, composite programs that use different effectful operations.

4.3. Example: square root

With these definitions in place, we can now formalize a refinement derivation such as the one in Fig. 1. However, we do so with a program which calculates the square root of a number using binary search, taken from our previous work. This program is not particularly interesting algorithmically, but illustrates how to use all of the refinement rules we have seen so far.

We begin by defining the specification:

Definition sqrtSpec (P : Ptr) : PT nat :=
   [(**exists** t, find s P = Some t), ($v^2 \leqslant t < (v+1)^2$)]

The precondition requires that P is a valid pointer in our heap, pointing to some value t. The postcondition guarantees that after execution, the result v will be the square root of the original value t.

Next we can sketch a candidate sqrt program:

```
t ← read P
Q ← new Q (t + 1)
R ← new R 0
while (read R + 1 ≢ read Q)
  q ← read Q
  r ← read R
  let mid := (q + r) / 2
  if (t < mid²)
  then write Q mid
  else write R mid
r ← read R
return r
```

The program starts by creating two auxiliary variables R and Q, which will serve as the search bounds. The lower bound R is initially set to 0, while the upper bound Q is set to $(t+1)$. The **while** loop then halves (using integer division) the interval between these bounds on every iteration until the bounds are consecutive numbers. The loop also needs an invariant, which should imply the post-condition. We define this in Coq as:

Definition sqrtInv (Q R : nat) (t : nat) (s : **heap** nat) :=
   exists q, r.find s Q = Some q ∧ find s R = Some r ∧ $R \neq Q$ ∧
         $r^2 \leqslant n < q^2$.

This invariant can be read as follows: for any value q and r referenced by pointers Q and R, respectively, the original value t must be in between the square of q and r. Upon completion, the program returns the value referenced by R.

Proving this program satisfies its specification, amounts to proving the following sqrtCorrect lemma.

Lemma sqrtCorrect (P : Ptr) :
   sqrtSpec P $\sqsubseteq$ sqrt P.

To do so, we can unfold the definitions of sqrt, sqrtSpec, and the semantics we defined above, yielding a complex verification condition.

This form of post-hoc verification is very different from the interactive program calculation that we would like to perform. We need to write the final program *before* we start our proof. After unfolding definitions and β-reduction, the proof goal that we are left with is often large and unwieldy.

In the next section we will develop machinery to enable the interactive discovery of programs, rather than the mere transcription of an existing proof.

5. Interactive refinement

Although we can now take any pen-and-paper proof of refinement and verify this in Coq, we are not yet playing to the strengths of the *interactive* theorem prover that we have at hand. In this section, we will show how to develop lemmas and definitions on top of those we have seen so far that facilitate the interactive calculation of a program from its specification.

We start by defining a function that determines when a statement is executable, i.e., when there are no occurrences of the Spec constructor:

> **Fixpoint** isExecutable (*c* : WhileL *a*) : Prop

Rather than fixing the exact program upfront, we can now reformulate the correctness lemma of swap as follows:

> **Definition** deriveSqrt (*P* : Ptr) :
> { *c* : WhileL nat
>   & (sqrtSpec *P*) ⊑ *c*
>   & isExecutable *c* }.

The notation $\{x : A \ \& \ P\,x \ \& \ Q\,x\}$ in Coq is used to denote a dependent triple consisting of a witness $x : A$, a proof that x satisfies the property P and a proof that x satisfies the property Q.

To prove this lemma we need to provide an executable c : WhileL nat and a proof that sqrtSpec ⊑ c. This is a superficial change – we could now complete the proof by providing our sqrt program as the witness c and reuse our previous correctness lemma. Instead of doing this, however, we wish to explore how to reformulate typical refinement calculus laws to enable the interactive construction of a suitable program.

We begin by observing that the semantics of our WhileL language proceeds structurally over the various language constructs. As a result, we can define a lemma for every language construct, describing precisely when a refinement step introducing that construct is valid. This corresponds to unfolding the definition of our semantics, for every construct of the WhileL language. For example, the rule for the Write statement is formulated below.

> **Lemma** writeRefines (*w w′* : WhileL *a*) (*ptr* : Ptr) (*y* : *v*)
>   (*d* : pre (semantics *w*) ⊆ pre (semantics (Write *ptr y w′*)))
>   (*h* : **forall** (*s* : *S*) (*p* : pre (semantics *w*) *s*) (*x* : *a*),
>   post (semantics *w′*) (update *s ptr y*) (snd (*d s p*)) *x*
>   ⊆ post (semantics *w*) *s p x*)
>   : *w* ⊑ Write *ptr y w′*.

Despite the apparent complexity, this lemma is trivial to prove: the two hypotheses are exactly what is needed to show that *w* is refined by Write *ptr v w′*. The first proof obligation states that the precondition of *w* should imply the precondition of Write *ptr v w′*; the second states that when the postcondition of *w′* holds on the *updated state*, i.e. after assigning *v* to the pointer *ptr*, the postcondition of *w* must also hold.

During the interactive refinement of a specification, however, we typically are not interested in refining two arbitrary WhileL programs, but rather calculating an executable WhileL program from its specification. Therefore, we can specialize the writeRefines lemma above to the case when the first program is a specification. Unfolding the definition of semantics and massaging the required hypotheses slightly yields the following lemma:

> **Lemma** writeSpec (*ptr* : Ptr) (*y* : *v*) (*spec* : PT *a*) (*w* : WhileL *a*)
>   (*H* : **forall** *s*, pre *spec s* → { *x* : *v* & find *s ptr* = Some *x* })
>   (*Step* : Spec
>   ([**fun** *s* ⇒ { *t* : *S* & prod (pre *spec t*) (*s* = (update *t ptr y*)) },
>   **fun** *s pres x s′* ⇒
>   (post *spec* (projT1 *pres*) (fst (projT2 *pres*)) *x s′*)]) ⊑ *w*) :
>   Spec *spec* ⊑ Write *b ptr y w*.

Essentially, this lemma states that to refine a specification with a Write statement, we need to prove that the pointer is already allocated on the heap (the argument *H*). Furthermore, we need to show after performing this update, we can refine the specification on the updated heap to some program *w*. Here the usage of relations to represent our postconditions introduce a bit of clutter, having to take deconstruct various parts of the precondition using the projections projT1, fst, and projT2 to formulate the desired postcondition.

We can define similar lemmas describing the verification conditions associated with introducing other constructors of the WhileL language, such as reading from memory, allocating a fresh pointer, introducing an if clause, or introducing a loop.

Note that in some cases, the continuation argument requires a more complex, higher-order hypothesis. For example, the readSpec lemma has the following general shape:

> **Lemma** readSpec (*ptr* : Ptr) (*spec* : PT *a*) (*w* : *b* → WhileL *a*)
>   (*H* : ...)
>   (*Step* : **forall** *v*, Spec [...]) ⊑ *w v*) :
>   Spec *spec* ⊑ Read *b ptr w*.

Here the *Step* argument quantifies over the result of the read, *v*; the specification makes precise that *v* is stored on the heap at address *ptr*. This complexity arises as we *are* interested in the result produced by Read, whereas Write does not produce and interesting result for the remainder of the computation. This is a natural generalization of the case for Write to handle those cases where some command produces a result used by the remaining computation.

We can now introduce custom tactics for each of these lemmas. These tactics apply the lemma and call some rudimentary proof automation:

> **Ltac** WRITE *ptr v* := eapply (writeSpec *ptr v*); simpl_goal.

Here the simpl_goal tactic unfolds various definitions, triggers β-reduction, and generally cleans up the proof context and goal. The tactic for While follows a similar structure:

> **Ltac** WHILE *I c* := eapply (whileSpec *I c*); simpl_goal.

The constructors that have a non-trivial continuation may introduce new variables, hence we adapt our tactics accordingly.

> **Ltac** READ *ptr v* := eapply (readSpec *ptr*); [| intros *v*]; simpl_goal.

We can also provide custom machinery for constructs defined in Coq itself, such as conditionals:

> **Lemma** ifSpec (*cond* : bool) (*spec* : PT *v a*) (*wt we* : WhileL *v a*) :
> (**if** *cond* **then** Spec *spec* $\sqsubseteq$ *wt* **else** Spec *spec* $\sqsubseteq$ *we*) $\rightarrow$
> (Spec *spec* $\sqsubseteq$ (**if** *cond* **then** *wt* **else** *we*)).

The associated tactic applies this lemma, remembers the condition, and splits the goal into two sub-goals accounting for the condition's value:

> **Ltac** IFF *c* := eapply (ifSpec *c*); remember *c* as *b*; destruct *b*; simpl_goal.

Finally, we can conclude any refinement calculation provided we have a suitable value and can show that the precondition implies the postcondition for any state *s*:

> **Lemma** returnStep (*v* : *a*) (*w* : WhileL *a*)
> (*H* : **forall** (*s* : *S*) (*P* : pre (semantics *w*) *s*), post (semantics *w*) *s* pre *v s*) :
> *w* $\sqsubseteq$ Return *v*.

We also define a corresponding tactic, RETURN *v*.

We now have the ingredients to tackle our problem: interactively proving deriveSqrt. That is, we must produce some code *c*, together with a proof that it refines our specification. The proof proceeds by introducing a new existential variable for *c*, which will be constructed while we perform the refinement calculation. We proceed by finding a refinement derivation by applying tactics one by one, that each modify the remaining specification. We have sketched the proof in Fig. 3. Although it leaves out several proof obligations, it is hopefully clear that the proof structure closely follows the code. Once we have shown our derivation correct, we still need to prove that the resulting program is executable—but this is trivial for any complete program.

The missing proof steps are not too long, but they detract from the overall structure of the proof that we wish to illustrate here. We refer the interested reader to our Coq development. Upon completing this proof, we can project out the *c* : WhileL unit value to obtain the verified program.

In this style, it is possible to start from a specification and incrementally write a program that satisfies it. Every step of the way, you can inspect how the specification evolves after applying a specific command. For larger developments, however, it may sometimes be desirable to intersperse verification and program development. In the calculation of sqrt above, the only real verification is done after each RETURN statement. One way to remedy this is by explicitly reformulating the specification:

> **Lemma** changeSpec (*pt*$_1$ *pt*$_2$: PT *a*) (*w* : WhileL *a*)
> (*d* : pre *pt*$_2$ $\subseteq$ pre *pt*$_1$)
> (*h* : **forall** (*s* : *S*) (*x* : pre *pt*$_2$ *s*) *v*, post *pt*$_1$ *s* (*d s x*) *v* $\subseteq$ post *pt*$_2$ *s x v*)
> (*H* : Spec *pt*$_1$ $\sqsubseteq$ *w*) :
> Spec *pt*$_2$ $\sqsubseteq$ *w*.

The changeSpec lemma follows immediately from transitivity of our refinement relation. Essentially, it allows us to replace the current specification with any other specification that refines it. We can also introduce a corresponding tactic:

> **Ltac** ASSERT *P* := unshelve eapply (changeSpec *P*).

```
econstructor.split.
− READ P t.
   NEW (s + 1) Q.
   NEW 0 R.
   WHILE (sqrtInv Q R s) (Cond Q R).
   + ... (* A proof of pre sqrtSpec → sqrtInv *)
   + READ Q vInQ.
     READ R vInR.
     IF (s <? ((vInQ + vInR) / 2)²)
     ∗ WRITE Q ((vInQ + vInR) / 2)
       RETURN tt.
       ... (* heap manipulation *)
     ∗ WRITE P ((vInQ + vInR) / 2)
       RETURN tt.
       ... (* heap manipulation *)
   + READ R vInR.
     RETURN tt.
     ... (* heap manipulation *)
− ... (* proving c to be executable *)
```

Fig. 3. Proof sketch of deriveSqrt.

This tactic applies the changeSpec lemma, thereby replacing the current specification with P. The unshelve tactic requires the user to complete the proof obligations associated with this application of changeSpec, before continuing the rest of the program calculation. Using these tactics, we are free to mix verification and calculational steps.

6. Case study: persistent arrays

As a final case study, we describe the derivation of a library for *persistent arrays*. A persistent array is a data-structure that stores a regular array and maintain the state of previous versions of the same array. Persistent arrays can be used as efficient building blocks of other data-structures, as in the union-find data-structure. In our development, we will follow the formalization by Conchon and Filliâtre [9] done previously in Coq.

The formalization by Conchon and Filliâtre uses two different modules. The first provides an implementation and interface for persistent arrays; whereas the second implements the union-find algorithm using this module. Any side effects are encapsulated in the persistent array module, providing a pure interface for the union-find structure. We will sketch how to calculate the persistent array module on which the union-find data structure relies using our library.

The module for persistent arrays has the following signature in OCaml:

```
module type PersistentArray = sig
  type a t
  val init : int → (int → a) → a t
  val get : a t → int → a
  val set : a t → int → a → a t
end
```

This module provides quite standard operations for persistent arrays. The function init takes an integer corresponding to the size of the array and functional representation of the array and returns a new persistent array t holding values of type a. The function get should takes such a value of type $a\ t$, an index of the array and returns the value stored there. Finally the function set accepts an array, an index and a value, and returns a new container with the old value on the given index replaced by the value passed as argument.

The derivation of each of these function will consist of three parts: a specification and an interactive proof demonstrating that we can find an implementation that meets the specification. For the sake of presentation, however, we will begin discussing the functions to familiarize readers with the algorithms we intend to derive. The code for these functions is given in Fig. 4.

The inductive data-type PAData represents the data to be stored in the heap. The Arr constructor holds an array, again represented as a function from indexes to values – for simplicity, we restrict arrays to only store natural numbers and we let this function be total. Just as in the work by Conchon and Filliâtre, the regular array lookup is represented as function application; whereas updating an array at some position with some value, can be done using the update operation.

The Diff constructor holds a position, a value, and a pointer to another array. It represents a single change in the *referenced* array at the given position with the given value. The Diff constructor is what makes the arrays *persistent*.

Inductive PAData : Type :=
   Arr (nat → nat)
   | Diff nat nat Ptr

init n f = new (Arr f) (**fun** ptr ⇒ Return ptr)

get t i =
   read t ⟫= **fun** $tval$ ⇒
   match $tval$ **with**
     | Arr a ⇒ lookup a i
     | Diff (j, v, t') ⇒
       if $i ≡ j$ **then** v **else** get t' i

set t i v =
   read t ⟫= **fun** $tval$ ⇒
   match $tval$ **with**
     | Arr a ⇒
      let old = lookup a i **in**
      new (Arr (update a i v)) ⟫= **fun** res ⇒
      write t (Diff i old res) ≫
      return res
     | Diff _ ⇒ new (Diff i v t)

Fig. 4. Functions manipulating persistent arrays.

The implementation of init simply creates a new array and returns the pointer associated to it. The get function recursively looks up a value in the chain of Diffs until it finds an Arr constructor. The set function behaves differently depending on the value referenced by t. In the Arr case, the array is modified, and t is set as a new indirection step and the (modified) array is returned. On the other hand, if the value in t is already a Diff, then a new indirection step is created an returned.

Note that Conchon and Filliâtre show a slightly different and more efficient way to implement get and set. However, this requires a new *reroot* function, which would introduce more complexity to our formalization.

We can now formalize specifications of the persistent array functions. When doing so, we notice that the important property to be verified is that each function should ensure that all pointers it uses should refer to a well-formed structure of a persistent array. In order to express that property, we define the following predicate, taken from [9]:

Inductive pa_model (s : **heap**) : Ptr → (nat → nat) → Type :=
   | pa_model_array :
     forall p f, find s p = Some (Arr f) → pa_model s p f
   | pa_mode_diff :
     forall p i v p', find s p = Some (Diff i v p') →
       forall f, pa_model s p' f →
     pa_model s p (update f i v).

This predicate allows us to relate a heap, a pointer and a function that represents the referenced array. One may informally explain this predicate as: pa_model s p f holds when a pointer p present in heap s, points to a valid persistant array represented by f. The constructor pa_model_array states that any pointer that references an array is a valid pointer and the function it represents is the f to which it is dereferenced. The constructor pa_mode_diff says that any pointer p referencing a Diff i v p' is also a valid pointer, as long as one can prove this property inductively for p'. The function representing the array will be an update of the resulting f in position i of v.

After having defined this predicate, we are now ready to formulate our specifications and derive their implementation.

6.1. The derivation of init

The init function is not only straightforward in its definition but also in its specification:

Definition initSpec (n : nat) (f : nat → nat) : PT Ptr :=
   [**fun** s ⇒ True
   , **fun** s $pres$ v s' ⇒
     (**forall** p' x : find s p' = Some x → $p' ≠ v$) ∧ pa_model s' v f].

The precondition is trivial: init may be called regardless of the current state of the heap. The postcondition is slightly more interesting. It states that the returned pointer v should be fresh and that it points to a valid array, represented by f. Deriving a suitable implementation for such the init function is straightforward.

6.2. The derivation of set

The specification for the set function is slightly more elaborate:

> **Definition** setSpec (ptr : Ptr) (i : nat) (v : nat) : PT Ptr :=
> [**fun** s ⇒ **exists** f, pa_model s ptr f
> , **fun** s $pres$ $newPtr$ s' ⇒
>   pa_model s' $newPtr$ (update f i v) ∧ pa_model s' ptr f].

The precondition requires a function f to exist, such that ptr points to a valid persistent array modeled by f. The post-condition ensures that, both the returned pointer and the original pointer will point to valid persistent arrays, and that the former is a suitably modified version the latter.

The derivation of the implementation of setSpec requires more manual proofs and requires several auxiliary properties relating the heap and pa_model. However, the application of these properties was quite systematic. Notably, this shows that there may be further room for customizable automation. When dealing with *custom* predicates such as pa_model, our rudimentary automation to simplify terms containing heaps does not have much effect. We believe that further effort, such as constructing suitable hint databases or tactic customization opportunities, the proof-to-program ratio can be much improved.

6.3. The derivation of get

The specification for the get function is straightforward:

> **Definition** getSpec (ptr : Ptr) (i : nat) : PT nat :=
> [**fun** s ⇒ **exists** f : pa_model s ptr f
> , **fun** s $pres$ v s' ⇒ $v = f$ i].

Somewhat surprisingly, this turned out to be the hardest of the three functions to derive. Despite the simplicity of the specification, the implementation is quite complex. Notably, it follows pointers in the heap and is not (obviously) structurally recursive. We chose to derive a variant of the original definition, introducing a While loop. The proof obligations can grow quite unwieldy – once again highlighting the need for further automation to keep their complexity in check.

We have sketched how to derive a library for persistent arrays. We expect that this development can be further combined with the development of other libraries that use persistent arrays as its key components, such as the existing Coq verification of the union-find data structure.

7. Discussion

The choice of our PT types and definition of refinement relation are not novel. Similar definitions of *indexed containers* [1] and *interaction structures* [14,15] can already be found in the literature. Indeed, part of this work was triggered by Peter Hancock's remark that these structures are closely related to *predicate transformers* and the refinement relation between them, as we have made explicit in this paper.

We are certainly not the first to explore the possibility of embedding a refinement calculus in a proof assistant. One of the first attempts to do so, to the best of our knowledge, was by Back and Von Wright [2]. They describe a formalization of several notions, such as weakest precondition semantics and the refinement relation, in the interactive theorem prover HOL. This was later extended to the *Refinement Calculator* [7], that built a new GUI on top of HOL using Tcl/Tk. More recently, Dongol et al. have extended these ideas even further in HOL, adding a separation logic and its associated algebraic structure [11]. There are far fewer such implementations in Coq, Boulmé [6] being one of the few exceptions. In contrast to the approach taken here, Boulmé explores the possibility of a monadic, shallow embedding, by defining the *Dijkstra Specification Monad*. Where Boulmé's work explores the lattice theoretic structure and fixpoint theory of refinement relation in Coq, it lacks custom refinement such as those presented here.

There is a great deal of work marrying effects and dependent types. Swierstra's thesis explores one potential avenue: defining a functional semantics for effects [21,23]. For some effects, such as non-termination, defining such a functional semantics in a total language is highly non-trivial. Therefore, systems such as Ynot take a different approach [19]. Ynot extends Coq with several axioms, corresponding to the different operations various effects support, such as reading from and writing to mutable state. The type of these axioms captures all the information that a programmer may use to reason about such effects. These types are similar to those presented here in Fig. 2. Contrary to the approach taken here, however, Ynot lets users write their programs without considering their specification. Users only need to write proofs after specifying

the pre- and postconditions for a certain function. The refinement calculus, on the other hand, starts from a specification, which is gradually refined to an executable program.

In the future, we hope to investigate how these various approaches to verification may be combined. One obvious next step would be to re-use the separation logic and associated proof automation defined by later installments of Ynot [8] as the model of the heap in our refinement calculus. Furthermore, we have (for now) chosen to ignore the variants associated with loops. As a result, the programs calculated may diverge. Embellishing our definitions with loop variants is straightforward, but will make our definitions even more cumbersome to use.

Type theory and the refinement calculus are both frameworks that combine specification and calculation. By embedding the refinement calculus in type theory, we study their relation further. The interactive structure of many proof assistants seems to fit well with the idea of *calculating* a program from its specification step-by-step. How well this approach scales, however, remains to be seen. For now, the embedding presented in this paper identifies an alternative point in the spectrum of available proof techniques for the construction of verified programs.

Acknowledgements

References

[1] Thorsten Altenkirch, Peter Morris, Indexed containers, in: 24th Annual IEEE Symposium on Logic In Computer Science, 2009, LICS'09, 2009.
[2] R.J.R. Back, J. von Wright, Refinement concepts formalized in higher order logic, Form. Asp. Comput. 2 (1989).
[3] Ralph-Johan Back, J. Von Wright, Refinement Calculus: A Systematic Introduction, Springer-Verlag Inc., New York, 1998.
[4] Ralph-Johan R. Back, Joakim von Wright, Refinement concepts formalised in higher order logic, Form. Asp. Comput. 2 (1) (1990) 247–272.
[5] R.J.R. Back, On the Correctness of Refinement in Program Development, Ph.D. thesis, University of Helsinki, 1978.
[6] Sylvain Boulmé, Intuitionistic refinement calculus, in: Typed Lambda Calculi and Applications, Springer, 2007, pp. 54–69.
[7] M.J. Butler, J. Grundy, T. Långbacka, R. Ruksenas, J. von Wright, The refinement calculator, in: Formal Methods Pacific, 1997.
[8] Adam Chlipala, Gregory Malecha, Greg Morrisett, Avraham Shinnar, Ryan Wisnesky, Effective interactive proofs for higher-order imperative programs, in: International Conference on Functional Programming, ICFP'09, 2009.
[9] Sylvain Conchon, Jean-Christophe Filliâtre, A persistent union-find data structure, in: Proceedings of the 2007 Workshop on ML, ACM, 2007, pp. 37–46.
[10] Edsger W. Dijkstra, A Discipline of Programming, Prentice–Hall, 1976.
[11] Brijesh Dongol, Victor B.F. Gomes, Georg Struth, A program construction and verification tool for separation logic, in: Mathematics of Program Construction, in: Lect. Notes Comput. Sci., vol. 9129, 2015.
[12] Robert W. Floyd, Assigning meanings to programs, Math. Asp. Comput. Sci. 19 (1967) 19–32.
[13] Peter Hancock, Pierre Hyvernat, Programming interfaces and basic topology, Ann. Pure Appl. Log. 137 (1) (2006) 189–239.
[14] Peter Hancock, Anton Setzer, Interactive programs in dependent type theory, in: Computer Science Logic, 2000, pp. 317–331.
[15] Peter Hancock, Anton Setzer, Specifying interactions with dependent types, in: Workshop on Subtyping and Dependent Types in Programming, 2000.
[16] Charles Antony Richard Hoare, An axiomatic basis for computer programming, Commun. ACM 12 (10) (1969) 576–580.
[17] Carroll Morgan, Programming from Specifications, Prentice–Hall, Inc., 1990.
[18] Aleksandar Nanevski, Greg Morrisett, Avi Shinnar, Paul Govereau, Lars Birkedal, Ynot: reasoning with the awkward squad, in: Proceedings of the Twelfth ACM SIGPLAN International Conference on Functional Programming, ICFP'08, 2008.
[19] Aleksandar Nanevski, Greg Morrisett, Avi Shinnar, Paul Govereau, Lars Birkedal, Ynot: dependent types for imperative programs, in: International Conference on Functional Programming, ICFP'08, 2008.
[20] Wouter Swierstra, A Hoare logic for the state monad, in: Theorem Proving in Higher Order Logics, Springer, 2009, pp. 440–451.
[21] Wouter Swierstra, A Functional Specification of Effects, Ph.D. thesis, University of Nottingham, 2009.
[22] Wouter Swierstra, Joao Alpuim, From proposition to program: embedding the refinement calculus in Coq, in: International Symposium on Functional and Logic Programming, Springer, 2016, pp. 29–44.
[23] Wouter Swierstra, Thorsten Altenkirch, Beauty in the beast: a functional semantics for the awkward squad, in: Haskell Workshop, 2007, pp. 25–36.