



A predicate transformer semantics for effects

Wouter Swierstra and Anne Baanen

Utrecht University

Constructive mathematics and computer programming†

BY P. MARTIN-LÖF

Department of Mathematics, University of Stockholm, Box 6701, S-113 85 Stockholm, Sweden

If programming is understood not as the writing of instructions for this or that computing machine but as the design of methods of computation that it is the computer's duty to execute (a difference that Dijkstra has referred to as the difference between computer science and computing science), then it no longer seems possible to distinguish the discipline of programming from constructive mathematics. This explains why the intuitionistic theory of types (Martin-Löf 1975 In *Logic Colloquium 1973* (ed. H. E. Rose & J. C. Shepherdson), pp. 73–118. Amsterdam: North-Holland), which was originally developed as a symbolism for the precise codification of constructive mathematics, may equally well be viewed as a programming language. As such it provides a precise notation not only, like other programming languages, for the programs themselves but also for the tasks that the programs are supposed to perform. Moreover, the inference rules of the theory of types, which are again completely formal, appear as rules of correct program synthesis. Thus the correctness of a program written in the theory of types is proved formally at the same time as it is being synthesized.

The day was closed by P. Martin-Löf... But the 50 minutes were not enough to introduce an ignorant audience to intuitionistic type theory to the extent that it could follow a comparison with Scottery. He was a very sympathetic speaker and convinced at least me that something (possibly even of great conceptual elegance) was going on.

Can we give a *constructive* account of
Dijkstra's weakest precondition semantics
in Martin-Löf type theory?

A predicate transformer semantics for effects

- A **predicate** on type a is some value of type

$a \rightarrow \text{Set}$

A predicate transformer semantics for effects

- A **predicate** on type a is some value of type

$a \rightarrow \text{Set}$

- A **predicate transformer** maps predicates to predicates:

$(a \rightarrow \text{Set}) \rightarrow (b \rightarrow \text{Set})$

A predicate transformer semantics for effects

- A **predicate** on type a is some value of type

$a \rightarrow \text{Set}$

- A **predicate transformer** maps predicates to predicates:

$(a \rightarrow \text{Set}) \rightarrow (b \rightarrow \text{Set})$

- A **predicate transformer semantics** assigns a predicate transformer to

$\text{wp} : (a \rightarrow b) \rightarrow (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$

$\text{wp} = \dots$

A predicate transformer semantics for effects

- A **predicate** on type a is some value of type

$a \rightarrow \text{Set}$

- A **predicate transformer** maps predicates to predicates:

$(a \rightarrow \text{Set}) \rightarrow (b \rightarrow \text{Set})$

- A **predicate transformer semantics** assigns a predicate transformer to

$\text{wp} : (a \rightarrow b) \rightarrow (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$

$\text{wp} = \cdot$

A predicate transformer semantics for effects

- A **predicate** on type a is some value of type

$a \rightarrow \text{Set}$

- A **predicate transformer** maps predicates to predicates:

$(a \rightarrow \text{Set}) \rightarrow (b \rightarrow \text{Set})$

- A **predicate transformer semantics** assigns a predicate transformer to

$\text{wp} : (a \rightarrow b) \rightarrow (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$

$\text{wp} = \cdot$

Or more generally, using *dependent types*

$\text{wp} : ((x : a) \rightarrow b\ x) \rightarrow (\forall x \rightarrow b\ x \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$

Example: predicates

We can illustrate the general principle using a (trivial) example:

$$\text{wp} : (a \rightarrow b) \rightarrow (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$$
$$\text{wp } f \text{ P } x = \text{P } (f \ x)$$

If we have a function `double` : $\mathbb{N} \rightarrow \mathbb{N}$ and the predicate:

$$\text{gt17} : \mathbb{N} \rightarrow \text{Set}$$
$$\text{gt17 } x = x > 17$$

What is the (weakest) precondition that needs to hold in order for the result of `double` to satisfy `gt17` – that is `double` produces a number greater than 17?

Example: predicates

We can illustrate the general principle using a (trivial) example:

$$\text{wp} : (a \rightarrow b) \rightarrow (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$$
$$\text{wp } f \ P \ x = P (f \ x)$$

If we have a function `double` : $\mathbb{N} \rightarrow \mathbb{N}$ and the predicate:

$$\text{gt17} : \mathbb{N} \rightarrow \text{Set}$$
$$\text{gt17 } x = x > 17$$

What is the (weakest) precondition that needs to hold in order for the result of `double` to satisfy `gt17` – that is `double` produces a number greater than 17?

$$Q : \mathbb{N} \rightarrow \text{Set}$$
$$Q \ x = \text{gt17 } (\text{double } x)$$

Example: relations

But many specifications *relate* inputs and outputs – instead of just requiring a number greater than 17, we may want a sorted permutation of our input list.

This follows naturally if you use dependent types.

Example: relations

But many specifications *relate* inputs and outputs – instead of just requiring a number greater than 17, we may want a sorted permutation of our input list.

This follows naturally if you use dependent types.

$\text{wp} : ((x : a) \rightarrow b\ x) \rightarrow (\forall x \rightarrow b\ x \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$

Consider the following (slightly contrived) example:

- take for $p : (xs : \text{List } a) \rightarrow \text{Permutation } xs$ as the first argument;
- and $\text{isSorted} : (xs : \text{List } a) \rightarrow \text{Permutation } xs \rightarrow \text{Set}$ as the second.

Then wp computes the precondition necessary for p to be a sorting function.

Computing with effects

So far this is not particularly exciting – it is no surprise that we can compute with predicates in Agda to reason about total functions.

But we can use the same techniques to reason about *effectful* functions.

Computing with effects

So far this is not particularly exciting – it is no surprise that we can compute with predicates in Agda to reason about total functions.

But we can use the same techniques to reason about *effectful* functions.

This is problematic at first glance – many proof assistants based on type theory are careful to *avoid* effects, but restrict themselves to a total language to ensure the soundness of the underlying logic.

Computing with effects

So far this is not particularly exciting – it is no surprise that we can compute with predicates in Agda to reason about total functions.

But we can use the same techniques to reason about *effectful* functions.

This is problematic at first glance – many proof assistants based on type theory are careful to *avoid* effects, but restrict themselves to a total language to ensure the soundness of the underlying logic.

For example, if we are not careful about handling unbounded recursion, we can define ‘bogus’ proofs such as:

```
silly :  $\forall x \rightarrow x < x$ 
```

```
silly x = silly x
```


Effects

Inspired by work on algebraic effects, we are careful separate **syntax** and **semantics**.

- A free monad fixes the syntax;
- the semantics is defined by a predicate transformer.

Our ICFP paper describes the syntax and semantics for a variety of different effects in this style:

- exceptions
- mutable state
- non determinism
- general recursion

Free monads

```
data Free (C : Set) (R : C → Set) (a : Set) : Set where
```

```
Pure : a → Free C R a
```

```
Step : (c : C) → (R c → Free C R a) → Free C R a
```

- A set C of *commands*;
- A function $R : C \rightarrow \text{Set}$ of responses associated with every command.

Different choices of C and R give arise to different effects.

Free monads

```
data Free (C : Set) (R : C → Set) (a : Set) : Set where
```

```
Pure : a → Free C R a
```

```
Step : (c : C) → (R c → Free C R a) → Free C R a
```

- A set C of *commands*;
- A function $R : C \rightarrow \text{Set}$ of responses associated with every command.

Different choices of C and R give rise to different effects.

- For example, to represent the familiar operations from the state monad, we can choose:
 - Commands $\text{Get} : C$ and $\text{Put} : s \rightarrow C$
 - Responses s for Get and \top for Put

State explicitly

Instantiating C and R accordingly yields the following data type (for some type of states $s : \text{Set}$):

```
data FS (a : Set) : Set where
```

```
  Get : (s → FS a) → FS a
```

```
  Put : s → FS a → FS a
```

```
  Return : a → FS a
```

If we choose s to be the natural numbers, we can write simple programs in this style:

```
incr : FS a
```

```
incr = get >>= λ x → put (x + 1) >> return x
```

Free monads: other examples

- Exceptions
 - Commands $\text{Abort} : C$
 - Responses \perp

Free monads: other examples

- Exceptions
 - Commands `Abort` : `C`
 - Responses \perp
- Non-determinism
 - Commands `Choice` : `C` and `Fail` : `C`
 - Responses `Bool` for `Choice` and \perp for `Fail`

Free monads: other examples

- Exceptions
 - Commands `Abort : C`
 - Responses \perp
- Non-determinism
 - Commands `Choice : C` and `Fail : C`
 - Responses `Bool` for `Choice` and \perp for `Fail`
- General recursion on a function $I \rightarrow O$
 - Commands `call : I \rightarrow C`
 - Responses `O`

Semantics for effects

In general, we want to study the meaning of Kleisli arrows – that is, programs of the form:

$a \rightarrow \text{Free } C \text{ R } b$

These correspond to ‘effectful programs’, taking an input of type a , performing effects from C and computing a value of type b .

Semantics for effects

In general, we want to study the meaning of Kleisli arrows – that is, programs of the form:

$$a \rightarrow \text{Free } C \text{ R } b$$

These correspond to ‘effectful programs’, taking an input of type a , performing effects from C and computing a value of type b .

Given our w_p function, we compute the weakest precondition associated with a Kleisli arrow:

$$w_p : (a \rightarrow \text{Free } C \text{ R } b) \rightarrow (\text{Free } C \text{ R } b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$$

But the postcondition here is expressed as a predicate on a free monad.

What happened to keeping syntax and semantics separate?

Semantics for effects

We'd like to define semantics with the following type:

$$(a \rightarrow \text{Free } C \text{ R } b) \rightarrow (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$$

But our wp semantics has the following form:

$$(a \rightarrow \text{Free } C \text{ R } b) \rightarrow (\text{Free } C \text{ R } b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$$

To do so, requires a predicate transformer semantics for effects:

$$(b \rightarrow \text{Set}) \rightarrow (\text{Free } C \text{ R } b \rightarrow \text{Set})$$

Defining predicate transformer semantics for effects boils down to defining such a function.

Semantics for effects – exceptions

$\text{wpPartial} : (a \rightarrow \text{Partial } b) \rightarrow (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$

$\text{wpPartial } f \ P = \text{wp } f \ (\text{mustPT } P)$

where

$\text{mustPT} : (b \rightarrow \text{Set}) \rightarrow (\text{Partial } b \rightarrow \text{Set})$

$\text{mustPT } P \ (\text{Pure } y) = P \ y$

$\text{mustPT } P \ (\text{Step } \text{Abort }) = \perp$

Here `Partial` refers to the free monad with a single command, `Abort`.

This semantics produces preconditions that guarantee `Abort` never happens.

Semantics for effects – exceptions

$\text{wpPartial} : (a \rightarrow \text{Partial } b) \rightarrow (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$

$\text{wpPartial } f \ P = \text{wp } f \ (\text{mustPT } P)$

where

$\text{mustPT} : (b \rightarrow \text{Set}) \rightarrow (\text{Partial } b \rightarrow \text{Set})$

$\text{mustPT } P \ (\text{Pure } y) = P \ y$

$\text{mustPT } P \ (\text{Step } \text{Abort }) = \perp$

Here `Partial` refers to the free monad with a single command, `Abort`.

This semantics produces preconditions that guarantee `Abort` never happens.

But other choices exist!

- Replace \perp with \top
- Require that P holds for some default value $d : a$
- ...

Semantics for effects – non-determinism

$\text{allPT} : (P : b \rightarrow \text{Set}) \rightarrow (\text{ND } b \rightarrow \text{Set})$

$\text{allPT } P \text{ (Pure } x) = P \ x$

$\text{allPT } P \text{ (Step Fail } k) = \top$

$\text{allPT } P \text{ (Step Choice } k) = \text{allPT } P \text{ (} k \text{ True)} \wedge \text{allPT } P \text{ (} k \text{ False)}$

Here we require P to hold for every possible result.

Semantics for effects – non-determinism

$\text{allPT} : (P : b \rightarrow \text{Set}) \rightarrow (\text{ND } b \rightarrow \text{Set})$

$\text{allPT } P \text{ (Pure } x) = P \ x$

$\text{allPT } P \text{ (Step Fail } k) = \top$

$\text{allPT } P \text{ (Step Choice } k) = \text{allPT } P \text{ (} k \text{ True)} \wedge \text{allPT } P \text{ (} k \text{ False)}$

Here we require P to hold for every possible result.

But again, alternatives exist.

The *gambler's nondeterminism* replaces \top with \perp and \wedge with \vee

Semantics for effects – state

$\text{statePT} : (P : (b \times s) \rightarrow \text{Set}) \rightarrow \text{FS } b \rightarrow (s \rightarrow \text{Set})$

$\text{statePT } P \text{ (Return } x) s = P (x, s)$

$\text{statePT } P \text{ (Get } k) s = \text{statePT } P (k s) s$

$\text{statePT } P \text{ (Put } s' k) s = \text{statePT } P k s'$

If necessary, we can also define a variant that takes an argument predicate:

$s \rightarrow (b \times s) \rightarrow \text{Set}$

So that we can observe the *relation* between input and output states.

State – larger example

```
data Tree (a : Set) : Set where
```

```
  Leaf : a → Tree a
```

```
  Node : Tree a → Tree a → Tree a
```

Exercise

Relabel such a binary tree with unique numbers assigned to each leaf.

State – larger example

```
data Tree (a : Set) : Set where  
  Leaf : a → Tree a  
  Node : Tree a → Tree a → Tree a
```

Exercise

Relabel such a binary tree with unique numbers assigned to each leaf.

```
relabel : Tree a → FS (Tree ℕ)  
relabel (Leaf _) = incr >>= Leaf  
relabel (Node l r) = relabel l >>= λ l' →  
                    relabel r >>= λ r' →  
                    return (Node l' r')
```

How do we show this is correct? Well to start with, we need a specification.

One way to specify the desired behaviour of our relabelling is:

$$P : \text{Tree } a \times \text{Nat} \rightarrow \text{Tree } \text{Nat} \times \text{Nat} \rightarrow \text{Set}$$
$$P (t, s) (t', s') = \text{flatten } t' \equiv \text{seq } s \ (\text{size } t)$$

Where $\text{seq } s \ x$ is the sequence of natural numbers starting from s of length x – it's easy to show that this does not contain duplicates.

One way to specify the desired behaviour of our relabelling is:

$$P : \text{Tree } a \times \text{Nat} \rightarrow \text{Tree } \text{Nat} \times \text{Nat} \rightarrow \text{Set}$$
$$P (t, s) (t', s') = \text{flatten } t' \equiv \text{seq } s \ (\text{size } t)$$

Where $\text{seq } s \ x$ is the sequence of natural numbers starting from s of length x – it's easy to show that this does not contain duplicates.

Unfortunately a direct proof showing that `relabel` satisfies this specification gets stuck quite quickly.

Compositionality

We do not yet know how to reason about composite programs written using binds.

But fortunately, we can prove a lemma along these lines:

```
compositionality : (c : FS a) (f : a → FS b) →  
  ∀ i P → statePT P (c >>= f) i ≡ statePT (wpState f P) c i
```

Compositionality

We do not yet know how to reason about composite programs written using binds.

But fortunately, we can prove a lemma along these lines:

```
compositionality : (c : FS a) (f : a → FS b) →  
  ∀ i P → statePT P (c >>= f) i ≡ statePT (wpState f P) c i
```

If you squint a bit, this is very similar to the usual relational composition used to reason about predicate transformers:

```
wp(c1 ; c2, R) = wp(c1, wp(c2, R))
```

Only here we have a monadic bind, passing an argument to f , rather the (more implicit) dependency between imperative programs.

Using this result, we can check that our relabelling function is indeed correct.

This shows how to assign a weakest precondition semantics to Kleisli arrows:

$$(a \rightarrow \text{Free } C \text{ R } b) \rightarrow (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$$

But why bother with such semantics in the first place?

It seems like a rather indirect way to reason about programs!

This shows how to assign a weakest precondition semantics to Kleisli arrows:

$$(a \rightarrow \text{Free } C \text{ R } b) \rightarrow (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$$

But why bother with such semantics in the first place?

It seems like a rather indirect way to reason about programs!

- We can also assign predicate transformer semantics to *specifications*;
- And use this semantics prove that a program satisfies its specification;
- Or even derive a program from its specification.

Specifications

Specifications

We define the following datatype of *specifications* on a function of type $(x : a) \rightarrow b\ x$

```
record Spec (a : Set) (b : a → Set) : Set where
```

```
field
```

```
  pre   : a → Set
```

```
  post  : (x : a) → b x → Set
```

- A *precondition* consisting of a predicate on a
- A *postcondition* consisting of a relation between $(x : a)$ and $b\ x$.

I'll often write such specifications as `[pre , post]`.

But how can we assign semantics to such specifications?

Semantics for specifications

$\text{wpSpec} : \text{Spec } a \ b \rightarrow (P : (x : a) \rightarrow b \ x \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})$

$\text{wpSpec } [\text{pre} , \text{post}] \ P = \lambda x \rightarrow (\text{pre } x) \wedge (\forall y \rightarrow \text{post } x \ y \rightarrow P \ x \ y)$

We can relate programs and specifications by relating the corresponding predicate transformers.

This idea – assigning predicate transformer semantics to *specifications* – is one of the key insights of the *refinement* calculus studied by Morgan, Back and von Wright.

Refinement

Given two predicate transformers, we can use the **refinement relation** to compare them:

$$\begin{aligned} _ \sqsubseteq _ &: (pt1\ pt2 : (b \rightarrow \text{Set}) \rightarrow (a \rightarrow \text{Set})) \rightarrow \text{Set} \\ pt1 \sqsubseteq pt2 &= \mathbf{forall}\ P\ x \rightarrow pt1\ P\ x \rightarrow pt2\ P\ x \end{aligned}$$

This relation is reflexitive, transitive and (morally) asymmetric.

Proving a program p satisfies it specification s amounts to showing:

$$wpSpec\ s \sqsubseteq wpEffect\ p$$

Refinements between programs

Not only can relate a program with its specification, but we can also compare two different programs using the refinement relation.

- For *pure* functions, $f \sqsubseteq g$ holds precisely when f and g are extensionally equal;
- For partial functions, $f \sqsubseteq g$ precisely when f and g agree on the domain of f ;
- For non-deterministic functions, $f \sqsubseteq g$ is equivalent to the subset relation.
- The gambler's non-deterministic semantics flips f and g .
- For state, $f \sqsubseteq g$ corresponds to the usual weaker-pres and stronger-posts.

Refinements between programs

Not only can relate a program with its specification, but we can also compare two different programs using the refinement relation.

- For *pure* functions, $f \sqsubseteq g$ holds precisely when f and g are extensionally equal;
- For partial functions, $f \sqsubseteq g$ precisely when f and g agree on the domain of f ;
- For non-deterministic functions, $f \sqsubseteq g$ is equivalent to the subset relation.
- The gambler's non-deterministic semantics flips f and g .
- For state, $f \sqsubseteq g$ corresponds to the usual weaker-pres and stronger-posts.

This is rather a nice result – the refinement relation captures the expected relation between effectful programs in a general way.

Compositionality of refinement

Pure functional programmers are spoiled. We're used to referential transparency, which allows us to employ equational reasoning.

Compositionality of refinement

Pure functional programmers are spoiled. We're used to referential transparency, which allows us to employ equational reasoning.

For all sensible predicate transformers, we can show the following result:

`compositionality : (f1 f2 : a → Free C R b) (g1 g2 : b → Free C R c) →`

`wp f1 ⊆ wp f2 →`

`wp g1 ⊆ wp g2 →`

`wp (f1 >=> g1) ⊆ wp (f2 >=> g2)`


Compositionality of refinement

Pure functional programmers are spoiled. We're used to referential transparency, which allows us to employ equational reasoning.

For all sensible predicate transformers, we can show the following result:

```
compositionality : (f1 f2 : a → Free C R b) (g1 g2 : b → Free C R c) →  
  wp f1 ⊆ wp f2 →  
  wp g1 ⊆ wp g2 →  
  wp (f1 >=> g1) ⊆ wp (f2 >=> g2)
```


Here we can reap the rewards of indirection: the verification of effectful programs is **compositional**.



In this fashion we can show a program—given by a Free C R a—satisfies some specification.

But can we **calculate** a program from its specification?

Program verification



In this fashion we can show a program—given by a $\text{Free } C \text{ R } a$ —satisfies some specification.

But can we **calculate** a program from its specification?


Let's consider values of the type $\text{Free } C \text{ R } (a + \text{Spec } a)$

We can assign them semantics by composing the semantics for specifications and effects.


[pre,post]




Program calculation



Program calculation




Program calculation



Typically, we prove little lemmas showing how each individual choice of command gives rise to new specifications, for which we must subsequently derive programs.

Program calculation



Typically, we prove little lemmas showing how each individual choice of command gives rise to new specifications, for which we must subsequently derive programs.

This style of calculation relies heavily on the **compositionality** of our semantics.

Even if you're not interested in program calculation, this gives you a 'small-step debugger' that you can use during *verification*.

All of our predicate transform semantics for effects have the following form:

$$\text{pt} : (b \rightarrow \text{Set}) \rightarrow (\text{Free } C \text{ R } b \rightarrow \text{Set})$$

But this is quite strange in a way – why is the predicate you return is *meaningful* in any way?

The degenerate case:

$$\text{pt } P \ c = \top$$

is type correct, but why is it still wrong?

Typically, we write these predicate transformers with an ‘intended’ semantics in mind.

`runState` : $FS\ a \rightarrow s \rightarrow a \times s$

We should show that the predicates we compute are *sound* with respect to these ‘handlers’.

In words, every result returned by this handler satisfies the desired postcondition when the computed precondition holds.

Typically, we write these predicate transformers with an ‘intended’ semantics in mind.

`runState : FS a → s → a × s`

We should show that the predicates we compute are *sound* with respect to these ‘handlers’.

In words, every result returned by this handler satisfies the desired postcondition when the computed precondition holds.

If you’re familiar with Dijkstra monads:

- the computational monad corresponds to this run function;
- the specification monad corresponds to the predicate transformer semantics.

More than a single effect?

So far, we've only talked about the semantics of different effects *in isolation*.

What about combining state and non-determinism? Or general recursion?

More than a single effect?

So far, we've only talked about the semantics of different effects *in isolation*.

What about combining state and non-determinism? Or general recursion?

- Free monads are closed under coproducts and compose nicely;
- Our predicate transformer semantics are defined as folds over free monads – these also compose nicely;
- We can put these together to study the predicate transformer semantics of compositions of effects.

Anne Baanen and I have a recent paper at MSFP where we use this to write parsers for regular languages.

Conclusion

- This gives a constructive & functional account of predicate transformer semantics.
- This approach works for a variety of different effects.
- We can relation effectful functions to their specifications in a compositional fashion.
- And even calculate programs from their spec.

Conclusion

- This gives a constructive & functional account of predicate transformer semantics.
- This approach works for a variety of different effects.
- We can relation effectful functions to their specifications in a compositional fashion.
- And even calculate programs from their spec.

Something (possibly of great conceptual elegance) is going on.