

Mixing Induction and Coinduction

Nils Anders Danielsson

University of Nottingham
nad@cs.nott.ac.uk

Thorsten Altenkirch

University of Nottingham
txa@cs.nott.ac.uk

Abstract

Purely inductive definitions give rise to tree-shaped values where all branches have finite depth, and purely coinductive definitions give rise to values where all branches are potentially infinite. If this is too restrictive, then an alternative is to use mixed induction and coinduction. This technique appears to be fairly unknown. The aim of this paper is to make the technique more widely known, and to present several new applications of it, including a parser combinator library which guarantees termination of parsing, and a method for combining coinductively defined inference systems with rules like transitivity.

The developments presented in the paper have been formalised and checked in Agda, a dependently typed programming language and proof assistant.

1. Introduction

Coinduction and corecursion are useful techniques for defining and reasoning about things which are potentially infinite. These “things” include streams and other (potentially) infinite data types (Coquand 1994; Giménez 1996; Turner 2004), subtyping relations for recursive types (Brandt and Henglein 1998; Gapeyev et al. 2002), process congruences (Milner 1990), congruences for functional programs (Gordon 1999), closures (Milner and Tofte 1991), semantics for divergence of programs (Cousot and Cousot 1992; Hughes and Moran 1995; Leroy and Grall 2009; Nakata and Uustalu 2009), and other applications.

However, the use of coinduction can lead to values which are “too infinite”. For instance, a non-trivial binary relation defined as a coinductive inference system cannot include the rule of transitivity, because a coinductive reading of transitivity would imply that every element is related to every other (to see this, build an infinite derivation consisting solely of uses of transitivity). As pointed out by Gapeyev et al. (2002) this is unfortunate, because without transitivity conceptually unrelated rules may have to be merged or otherwise modified, in order to ensure that transitivity can be proved as a derived property. Gapeyev et al. give the example of subtyping for records, where a dedicated rule of transitivity ensures that one can give separate rules for depth subtyping (which states that a record field type can be replaced by a subtype), width subtyping (which states that new fields can be added to a record), and permutation of record fields.

Fortunately this problem can be solved. The problem stems from a *coinductive* reading of transitivity, and it can be solved by reading the rule of transitivity *inductively*, and only using coinduction where it is necessary. In Section 4 this idea is illustrated by using mixed induction and coinduction to define a subtyping relation for recursive types. The rule which defines when a function type is a subtype of another is defined coinductively (following Brandt and Henglein (1998) and Gapeyev et al. (2002)), while the rule of transitivity is defined inductively.

It should be noted that some caution is necessary when using mixed induction and coinduction to define inference systems. Given an inductively defined inference system it is always sound to include an admissible property as a new inductive rule. However, in the presence of coinduction such an inclusion may not be sound if the admissible rule does not have a sufficiently “contractive” proof; see Section 4.5 for more details. (This problem does not affect our definition of subtyping, because we prove that the definition coincides with other definitions from the literature.)

The technique of mixing induction and coinduction is not limited to defining transitive inference systems. It is generally useful when defining sets of tree-shaped values (like the typing derivations above) where some values are guaranteed to have finite depth (typing derivations consisting solely of the rules of transitivity and reflexivity, for instance), and other values are allowed to have infinite depth.

As a second example of mixed induction and coinduction a library of parser combinators (Burge 1975; Wadler 1985; Fairbairn 1987, and many others) is defined. Parser combinators can provide an elegant and declarative method for specifying parsers. When compared with parser generators they have some advantages: it is easy to abstract over recurring grammatical patterns, and there is no need to use a separate tool just to parse something. On the other hand there are also some disadvantages: there is a risk of lack of efficiency, and parser generators can give static guarantees about termination and non-ambiguity which most parser combinator libraries fail to give. The library defined here addresses one of these points by ensuring statically that parsing will terminate. It also comes with a formal semantics.

When using parser combinators the parsers/grammars are often constructed using cyclic definitions, so it is natural to see the definitions as being partly corecursive. However, a purely coinductive reading of the choice and sequencing combinators would allow definitions like $p = p \mid p$ and $p' = p' \cdot p'$, for which it is impossible to implement parsing in a total way (in a pure setting where p and p' can only be inspected via their infinite unfoldings). As shown in Section 5 totality can be ensured by reading choice inductively, and only reading an argument of the sequencing operator coinductively when the other argument does not accept the empty string.

The examples above show how typing derivations and term languages can benefit from the use of mixed induction and coinduction. In Section 6 the technique is applied to the definition of a semantics: it is shown how a big-step semantics which handles both

converging and diverging computations can be defined without duplication of rules.

Guarded corecursion (see Section 2) is used throughout the paper to define values of coinductive types. Guardedness is only an approximation of productivity, which essentially states that every finite approximation of a conceptually infinite value should be computable in a finite number of steps. The restriction to guardedness can make programs and proofs harder to define, but is easy to check mechanically, and is used in some form in several programming languages/proof assistants. Section 3 presents a method for working around the limitations of guardedness. This method is admittedly ad-hoc, but is included because it uses mixed induction and coinduction, and because the technique is used to show productivity in other parts of the paper.

The method of mixing induction and coinduction is not new; Section 7.1 lists further applications of the technique. However, the method does not seem to be well-known. For instance, several authors mention that transitivity and coinduction cannot be combined (Gapeyev et al. 2002; Levin and Pierce 2003; Colazzo and Ghelli 2005), but as mentioned above this is not generally true. The main purpose of this paper is to make this technique more well-known.

To summarise, the main contributions of this paper are as follows:

- The technique of mixing induction and coinduction is explained, and its utility is demonstrated using several examples.
- It is shown that it is possible to combine coinduction and the rule of transitivity when defining an inference system.
- It is shown how parser combinators can be implemented in such a way that termination is guaranteed.
- A new method for working around the limitations of guardedness is explained.

Furthermore all the main developments have been formalised using the experimental, dependently typed, total¹ functional programming language Agda (Norell 2007; Agda Team 2009), and the source code can (at the time of writing) be downloaded from the first author's web page.

The rest of the paper is structured as follows: Section 2 gives an introduction to coinduction, and Section 3 describes a technique for working around the limitation of guardedness. Section 4 discusses subtyping for recursive types, Section 5 describes a parser combinator library which guarantees termination, and Section 6 discusses a big-step semantics which handles both converging and diverging computations without duplication of rules. Finally Section 7 discusses related work and Section 8 concludes.

2. Coinduction

Coinduction and corecursion can be defined in several ways. For concreteness this paper uses one particular method, but the main ideas should be generally valid (in theories which are sufficiently strong). In order to show that the techniques which are explained can be used in practical formalisations and programs, to reduce the risk of errors, and to make things concrete, the definitions in this paper will be carried out in Agda, slightly modified² to make things simpler and aid readability.

¹ Note that the meta-theory of Agda has not been properly formalised, so take statements such as “total” with a grain of salt. However, even if Agda should turn out to be unsound the definitions and proofs in the paper are still valid constructions in intuitionistic type theory.

² We have omitted some type declarations for implicit arguments, simplified the definitional equality of corecursive programs, and ignored universe levels ($Set_0, Set_1 \dots$).

This brief introduction does not attempt to explain coinduction in detail, but rather aims to give the intuition necessary to understand the rest of the paper. Coinductive definitions can be viewed as the categorical dual of inductive definitions (Hagino 1987). In the context of intuitionistic type theory coinductive types have been investigated by Mendler (1988), Coquand (1994), and Giménez (1996), and Leroy and Grall (2009) relate one type theoretic approach to coinduction with the approach based on the Knaster-Tarski fixpoint theorem.

The rest of this section illustrates the approach to coinduction taken in Agda. To start with, let us define the type of infinite streams:

```
data Stream (A : Set) : Set where
  _::_ : A → ∞ (Stream A) → Stream A
```

This states that *Stream* is a data type parametrised on the element type (or set) *A*, with a single constructor *_::_* which takes two arguments (*_::_* is an infix operator; the underscores mark the argument positions). The type function $\infty : Set \rightarrow Set$ marks an argument as being *coinductive*. Data type definitions can be seen as defining fixpoints of functors, and a data type definition $T = F (\infty T) T$, where *F* is a strictly positive functor in two arguments which does not mention *T*, should be read as the nested fixpoint $\nu C. \mu I. F C I$ (almost; see Section 7.1). Here $\mu X.G X$ and $\nu X.G X$ stand for the initial algebra and terminal coalgebra, respectively, of a given functor *G*. In the case of *Stream* we get the expected interpretation $Stream A = \nu C.A \times C$, modulo the presence of the named constructor *_::_*.

Programmers may prefer to think about ∞ as the suspension type constructor which is used to implement non-strictness in strict languages (Wadler et al. 1998). Just as the suspension type constructor the function ∞ comes with delay and force operators, here called $\#_-$ and b :

```
 $\#_- : \forall \{A\} \rightarrow A \rightarrow \infty A$ 
 $^b : \forall \{A\} \rightarrow \infty A \rightarrow A$ 
```

($\#_-$ is a tightly binding prefix operator; ordinary function application binds tighter, though. Arguments in braces, $\{ \dots \}$, are implicit, and do not need to be given explicitly as long as Agda can infer them from the context.)

Functions constructing coinductive values, like the *map* function for streams, can be defined by corecursion:

```
map : ∀ {A B} → (A → B) → Stream A → Stream B
map f (x :: xs) = f x :: # map f (^b xs)
```

Agda is a total language, so functions which construct infinite values have to be productive; even if the value being constructed is infinite it should always be possible to compute the next constructor in a finite number of steps. This is enforced by limiting corecursion to *guarded* corecursion (Coquand 1994). The definition of *map* is accepted because the corecursive call is guarded by the coinductive constructor $\#_-$, without any non-constructor function between the left-hand side and the corecursive call.

The *map* function is easy to define using guarded corecursion. However, sometimes guardedness is an inconvenient restriction. Fortunately guardedness is only an approximation of productivity, used because it is easy to check mechanically. When working informally, or in a formal setting where productivity is established semantically, guardedness can be replaced by more liberal principles.

The uses of $\#_-$ and b in the definition of *map* can be perceived as cluttering the code. One could imagine designing a language in which these operators can be inferred automatically, based on type

information. However, for the purposes of this paper it seems better to keep this information explicit.

Let us now see how coinduction can be mixed with induction. Hancock et al. (2009) define a language of stream processors, taking streams of A s to streams of B s, as the nested fixpoint $\nu C. \mu I. (A \rightarrow I) + B \times C$. This language can be represented using the following data type:

```
data SP (A B : Set) : Set where
  get  : (A → SP A B) → SP A B
  put  : B → ∞ (SP A B) → SP A B
```

(Note that $(A B : Set)$ means that both the parameters A and B have type Set ; it is not an application of A to B .) The recursive argument of `get` is inductive, while the recursive argument of `put` is coinductive. This means that a stream processor can only read a finite number of elements from the input before having to produce some output. The semantics of stream processors can be defined as follows:

```
[[_]] : ∀ {A B} → SP A B → Stream A → Stream B
[[ get f   ]] (a :: as) = [[ f a ]] (ᵇ as)
[[ put b sp ]] as      = b :: ᵇ [[ sp ]] as
```

In the case of `get` one element from the input stream is consumed, and potentially used to guide the rest of the computation, while in the case of `put` one output element is produced. The definition of `[[_]]` uses a lexicographic combination of guarded corecursion and structural recursion: in the second clause the corecursive call is guarded, while in the first clause the recursive call “preserves guardedness” (it takes place under zero coinductive constructors rather than at least one) and the stream processor argument is structurally smaller.

Let us now turn to coinductively defined *relations*. Equality of streams can be defined as follows:

```
data _≈_ {A : Set} : Stream A → Stream A → Set where
  _::_ : (x : A) {xs ys : ∞ (Stream A)} →
    ∞ (ᵇ xs ≈ ᵇ ys) → x :: xs ≈ x :: ys
```

This definition states that two streams are equal whenever their heads are (definitionally) equal and their tails are, recursively, equal. Note that this definition should be read coinductively, and that the tails are forced in the recursive argument’s type. Note also the use of the dependent function space $(x : A) \rightarrow B$, where B can depend on x . Note finally that constructors can be overloaded in Agda.

Elements of coinductively defined relations can be constructed by using corecursion. As an example, let us prove the *map-iterate* property. The function *iterate* repeatedly applies a function to a seed element and collects the results in a stream, $iterate\ f\ x = x :: \# (f\ x :: \# (f\ (f\ x) :: \dots))$:

```
iterate : ∀ {A} → (A → A) → A → Stream A
iterate f x = x :: ᵇ iterate f (f x)
```

The *map-iterate* property can then be proved by using guarded corecursion (the term *guarded coinduction* could also be used):

```
map-iterate : ∀ {A} (f : A → A) (x : A) →
  map f (iterate f x) ≈ iterate f (f x)
map-iterate f x = f x :: ᵇ map-iterate f (f x)
```

To see how this proof works, consider how it can be built up step by step (as in an interactive Agda session):

```
map-iterate f x = ?
```

The type of the *goal* ? is $map\ f\ (iterate\ f\ x) \approx iterate\ f\ (f\ x)$. Agda types should always be read up to normalisation, so this is equivalent to

```
f x :: ᵇ map f (ᵇ (ᵇ iterate f (f x))) ≈
f x :: ᵇ iterate f (f (f x)) .
```

(Note that normalisation does not involve evaluation under $\#$, and that $\# (ᵇ x)$ reduces to x .) This type matches the result type of the equality constructor `_::_`, so we can refine the goal:

```
map-iterate f x = f x :: ?
```

The new goal type is

```
∞ (map f (iterate f (f x)) ≈ iterate f (f (f x))),
```

so the proof can be finished by an application of the coinductive hypothesis under the guarding constructor $\#$:

```
map-iterate f x = f x :: ᵇ map-iterate f (f x)
```

3. An ad-hoc method for making corecursive definitions guarded

As mentioned in Section 2 guardedness is sometimes an inconvenient restriction: there are productive programs which are not syntactically guarded. This section illustrates a technique which can be used to work around this restriction.

As an example, consider the following definition of the stream of Fibonacci numbers:

```
fib : Stream ℕ
fib = 0 :: ᵇ zipWith _+_ fib (1 :: ᵇ fib)
```

The definition uses the function *zipWith*, which combines two streams:

```
zipWith : ∀ {A B C} → (A → B → C) →
  Stream A → Stream B → Stream C
zipWith f (x :: xs) (y :: ys) = f x y :: ᵇ zipWith f (ᵇ xs) (ᵇ ys)
```

While the definition of *fib* is productive, it is not guarded, because the function *zipWith* is not a constructor. If *zipWith* were a constructor the definition would be guarded, though, and this presents a way out: we can define a problem-specific language which includes *zipWith* as a constructor, and then define an interpreter for the language by using guarded corecursion.

A simple language of stream programs can be defined as follows:

```
data StreamP : Set → Set where
  _::_ : ∀ {A} → A → ∞ (StreamP A) → StreamP A
  zipWith : ∀ {A B C} → (A → B → C) →
    StreamP A → StreamP B → StreamP C
```

Note that the stream program argument of `_::_` is coinductive, while the arguments of `zipWith` are inductive; a stream program consisting of an infinitely deep application of `zipWith`s would not be productive.

Stream programs will be turned into streams in two steps. First a kind of weak head normal form (WHNF) for stream programs is computed recursively, and then the resulting stream is computed corecursively. The WHNFs are defined in the following way:

```
data StreamW : Set → Set where
  _::_ : ∀ {A} → A → StreamP A → StreamW A
```

Note that the stream argument to `_::_` is a (“suspended”) program, not a WHNF. The function *whnf* which computes WHNFs can be defined by structural recursion:

```

whnf : ∀ {A} → StreamP A → StreamW A
whnf (x :: xs) = x :: b xs
whnf (zipWith f xs ys) with whnf xs | whnf ys
whnf (zipWith f xs ys) | x :: xs' | y :: ys' =
  f x y :: zipWith f xs' ys'

```

(The **with** construct is used to pattern match on the results of intermediate computations.) WHNFs can then be turned into streams corecursively:

```

mutual
  [ ]W : ∀ {A} → StreamW A → Stream A
  [ x :: xs ]W = x :: # [ xs ]P
  [ ]P : ∀ {A} → StreamP A → Stream A
  [ xs ]P = [ whnf xs ]W

```

Note that this definition is guarded. (Agda accepts definitions like this one even though it is split up over two mutually defined functions.) Note also that this method can turn problems with productivity into problems with termination: if the definition of `zipWith` had not been productive the definition of `whnf` would not have been terminating.

Given the language above we can now define the stream of Fibonacci numbers using guarded corecursion:

```

fibP : StreamP ℕ
fibP = 0 :: # zipWith _+_ fibP (1 :: # fibP)
fib : Stream ℕ
fib = [ fibP ]P

```

One can prove that this definition satisfies the original equation for `fib` by first proving coinductively that `[]P` is homomorphic with respect to `zipWith`/`zipWith`:

```

zipWith-hom : ∀ {A B C} (f : A → B → C) xs ys →
  [ zipWith f xs ys ]P ≈ zipWith f [ xs ]P [ ys ]P
fib-correct : fib ≈ 0 :: # zipWith _+_ fib (1 :: # fib)

```

(The proofs are omitted.)

As stated in the introduction the method presented in this section is ad-hoc: when a new problem is encountered a new language such as `StreamP` may have to be defined. However, when working in a language which requires corecursion to be guarded we have found the method to be useful, and it readily generalises to other types than streams, and to developments using several types. In fact, the method is probably more useful for defining proofs (for instance subtyping proofs; see Sections 4.3 and 4.4) than for defining ordinary programs, because when defining a proof the computational overhead of the coding is usually irrelevant.

4. Subtyping for recursive types

This section shows how a coinductively defined inference system can be combined with the rule of transitivity. The technique is general, but illustrated in the context in which the problem was raised: subtyping for recursive types (Gapeyev et al. 2002).

4.1 Syntax

Brandt and Henglein (1998) define the following language of recursive types:

$$\sigma, \tau ::= \perp \mid \top \mid X \mid \sigma \rightarrow \tau \mid \mu X. \sigma \rightarrow \tau$$

Here \perp and \top are the least and greatest types, respectively, X is a variable, $\sigma \rightarrow \tau$ is a function type, and $\mu X. \sigma \rightarrow \tau$ is a fixpoint, with bound variable X . (The body of the fixpoint is required to be a function type, so types like $\mu X. X$ are ruled out.) The intention

is that a fixpoint $\mu X. \sigma \rightarrow \tau$ should be equivalent to its unfolding $(\sigma \rightarrow \tau)[X := \mu X. \sigma \rightarrow \tau]$.

The language above can be represented in Agda as follows:

```

data Ty (n : ℕ) : Set where
  ⊥       : Ty n
  ⊤       : Ty n
  var     : Fin n → Ty n → Ty n
  →→     : Ty n → Ty n → Ty n
  μ→→    : Ty (1 + n) → Ty (1 + n) → Ty n

```

Here variables are represented using de Bruijn indices: `Ty n` represents types with at most n free variables, and `Fin n` is a type representing the first n natural numbers. Substitution can also be defined; $\sigma [\tau]$ is the capture-avoiding substitution of τ for variable 0 in σ :

$$-[_] : \forall \{n\} \rightarrow Ty (1 + n) \rightarrow Ty n \rightarrow Ty n$$

The following function unfolds a fixpoint one step:

```

unfold(μ→→) : \{n\} \rightarrow Ty (1 + n) \rightarrow Ty (1 + n) \rightarrow Ty n
unfold(μ→→) = (\sigma \rightarrow \tau) [ \mu \sigma \rightarrow \tau ]

```

(Note that `-[_]` and `unfold(μ→→)` are both infix operators which take two arguments.)

4.2 Subtyping via trees

A natural definition of subtyping goes via subtyping for potentially infinite trees (Gapeyev et al. 2002):

```

data Tree (n : ℕ) : Set where
  ⊥       : Tree n
  ⊤       : Tree n
  var     : Fin n → Tree n → Tree n
  →→     : ∞ (Tree n) → ∞ (Tree n) → Tree n

```

The subtyping relation for trees can be given coinductively as follows:

```

data ≤Tree : Tree n → Tree n → Set where
  ⊥       : ⊥ ≤Tree τ
  ⊤       : σ ≤Tree ⊤
  var     : var x ≤Tree var x
  →→     : ∞ (τ1 ≤Tree σ1) → ∞ (τ2 ≤Tree σ2) →
    σ1 → σ2 ≤Tree τ1 → τ2

```

Note the contravariant treatment of the codomain of the function space. (Note also that the declarations of bound variables like n and τ have been omitted to reduce clutter.)

The semantics of a recursive type can be given in terms of its unfolding as a potentially infinite tree:

```

[ ] : \{n\} \rightarrow Ty n \rightarrow Tree n
[ ⊥ ] = ⊥
[ ⊤ ] = ⊤
[ var x ] = var x
[ σ \rightarrow τ ] = # [ σ ] \rightarrow # [ τ ]
[ μ \sigma \rightarrow \tau ] = # [ σ [ χ ] ] \rightarrow # [ τ [ χ ] ]
  where χ = μ \sigma \rightarrow \tau

```

The subtyping relation for types can then be defined by combining `≤Tree` and `[]`:

```

≤Type : \{n\} \rightarrow Ty n \rightarrow Ty n \rightarrow Set
σ ≤Type τ = [ σ ] ≤Tree [ τ ]

```

Amadio and Cardelli (1993) also define subtyping for recursive types by going via potentially infinite trees, but they define the subtyping relation *inductively* on finite trees, and define that an infinite tree σ is a subtype of another tree τ when every finite approximation (of a certain kind) of σ is a subtype of the correspond-

ing approximation of τ . It is easy to show that this definition, as adapted by Brandt and Henglein (1998), is equivalent to the one given above. One direction of the proof uses induction on the depth of approximation, and the other constructs elements of $\sigma \leq_{\text{Tree}} \tau$ corecursively (see the code which accompanies the paper).

4.3 Subtyping using mixed induction and coinduction

Subtyping can also be defined directly, without going via trees. The following definition is inspired by Brandt and Henglein (1998), see Section 4.4:

```

data  $\_ \leq \_ : \text{Ty } n \rightarrow \text{Ty } n \rightarrow \text{Set}$  where
   $\perp$       :  $\perp \leq \tau$ 
   $\top$       :  $\sigma \leq \top$ 
   $\_ \rightarrow \_$   :  $\infty (\tau_1 \leq \sigma_1) \rightarrow \infty (\sigma_2 \leq \tau_2) \rightarrow$ 
              $\sigma_1 \rightarrow \sigma_2 \leq \tau_1 \rightarrow \tau_2$ 
  unfold :  $\mu \tau_1 \rightarrow \tau_2 \leq \text{unfold}(\mu \tau_1 \rightarrow \tau_2)$ 
  fold   :  $\text{unfold}(\mu \tau_1 \rightarrow \tau_2) \leq \mu \tau_1 \rightarrow \tau_2$ 
  refl   :  $\tau \leq \tau$ 
  trans  :  $\tau_1 \leq \tau_2 \rightarrow \tau_2 \leq \tau_3 \rightarrow \tau_1 \leq \tau_3$ 

```

Note that the structural rules (\perp , \top , $_ \rightarrow _$) are defined coinductively, while the other rules, most importantly **trans**, are defined inductively. Hence, assuming that the relation expresses the right thing, this definition shows that coinduction and the rule of transitivity can be combined.

Showing that the relation expresses the right thing, i.e. that it is equivalent to the one given in Section 4.2 (and thus equivalent to Amadio and Cardelli's), is straightforward:

```

complete :  $\sigma \leq_{\text{Type}} \tau \rightarrow \sigma \leq \tau$ 
sound    :  $\sigma \leq \tau \rightarrow \sigma \leq_{\text{Type}} \tau$ 

```

Completeness can be proved by a simple application of guarded corecursion. The soundness proof is a little more tricky. Assume that we have proved that $_ \leq_{\text{Type}} _$ is transitive:

```

transType :  $\tau_1 \leq_{\text{Type}} \tau_2 \rightarrow \tau_2 \leq_{\text{Type}} \tau_3 \rightarrow \tau_1 \leq_{\text{Type}} \tau_3$ 

```

One might be tempted to prove soundness by a simple application of lexicographic guarded corecursion and structural recursion, the technique used to define $\llbracket _ \rrbracket$ in Section 2, but this fails. Consider the case for **trans**:

```

sound (trans  $\tau_1 \leq \tau_2$   $\tau_2 \leq \tau_3$ ) =
  transType (sound  $\tau_1 \leq \tau_2$ ) (sound  $\tau_2 \leq \tau_3$ )

```

(Note that $\tau_1 \leq \tau_2$ and $\tau_2 \leq \tau_3$ are variable names.) Here the recursive calls are on smaller arguments, but **trans**_{Type} is not a constructor, so guardedness is not preserved. If **trans**_{Type} were a constructor the proof would be accepted, though; the proof can be rescued by an application of the technique presented in Section 3.

4.4 Inductive axiomatisation of subtyping

Brandt and Henglein (1998) define subtyping essentially as in Section 4.3, but instead of using mixed induction and coinduction they encode coinduction inductively. Their subtyping relation is ternary: $A \vdash \sigma \leq \tau$ means that σ is a subtype of τ given the assumptions in A . An assumption (a hypothesis) is simply a pair of types:

```

data Hyp ( $n : \mathbb{N}$ ) : Set where
   $\_ , \_ : \text{Ty } n \rightarrow \text{Ty } n \rightarrow \text{Hyp } n$ 

```

The subtyping relation is defined as follows:

```

data  $\_ \vdash \_ \leq \_ (A : \text{List } (\text{Hyp } n)) :$ 
   $\text{Ty } n \rightarrow \text{Ty } n \rightarrow \text{Set}$  where
   $\perp$       :  $A \vdash \perp \leq \tau$ 
   $\top$       :  $A \vdash \sigma \leq \top$ 

```

```

 $\_ \rightarrow \_$  : let  $H = (\sigma_1 \rightarrow \sigma_2, \tau_1 \rightarrow \tau_2)$  in
   $H :: A \vdash \tau_1 \leq \sigma_1 \rightarrow H :: A \vdash \sigma_2 \leq \tau_2 \rightarrow$ 
   $A \vdash \sigma_1 \rightarrow \sigma_2 \leq \tau_1 \rightarrow \tau_2$ 
unfold :  $A \vdash \mu \tau_1 \rightarrow \tau_2 \leq \text{unfold}(\mu \tau_1 \rightarrow \tau_2)$ 
fold   :  $A \vdash \text{unfold}(\mu \tau_1 \rightarrow \tau_2) \leq \mu \tau_1 \rightarrow \tau_2$ 
refl   :  $A \vdash \tau \leq \tau$ 
trans  :  $A \vdash \tau_1 \leq \tau_2 \rightarrow A \vdash \tau_2 \leq \tau_3 \rightarrow A \vdash \tau_1 \leq \tau_3$ 
hyp    :  $(\sigma, \tau) \in A \rightarrow A \vdash \sigma \leq \tau$ 

```

Note that coinduction is encoded in the $_ \rightarrow _$ rule by inclusion of the consequent in the lists of assumptions of the antecedents.

Brandt and Henglein prove that their relation (with an empty list of assumptions) is equivalent to Amadio and Cardelli's. Their proof is considerably more complicated than the proof outlined above which shows that $_ \leq _$ is equivalent to Amadio and Cardelli's definition, but as part of the proof they show that subtyping is decidable. By composing the two equivalence proofs we get that subtyping as defined in Section 4.3 is also decidable. However, Brandt and Henglein use a classical argument to show that their algorithm terminates, so it is not entirely obvious that it can be implemented in a total, constructive type theory like Agda. Fortunately some small modifications suffice to show that subtyping is decidable in this setting; details are available in the code accompanying the paper.

It may be interesting to see how the inductive encoding of coinduction used in $_ \vdash _ \leq _$ can be turned directly into the "actual" coinduction used in $_ \leq _$, i.e. to see how soundness of $_ \vdash _ \leq _$ can be proved: it can be done using a cyclic proof. To state soundness the type *All* is used; *All P xs* means that all elements in *xs* satisfy *P*:

```

data All ( $P : A \rightarrow \text{Set}$ ) : List  $A \rightarrow \text{Set}$  where
   $[]$  : All  $P []$ 
   $\_ :: \_ : P x \rightarrow \text{All } P xs \rightarrow \text{All } P (x :: xs)$ 

```

The soundness proof shows that if $A \vdash \sigma \leq \tau$, where all pairs (σ', τ') in A satisfy $\sigma' \leq \tau'$, then $\sigma \leq \tau$:

```

Valid :  $(\text{Ty } n \rightarrow \text{Ty } n \rightarrow \text{Set}) \rightarrow \text{Hyp } n \rightarrow \text{Set}$ 
Valid  $\_ R \_ (\sigma_1, \sigma_2) = \sigma_1 R \sigma_2$ 
sound : All (Valid  $\_ \leq \_$ )  $A \rightarrow A \vdash \sigma \leq \tau \rightarrow \sigma \leq \tau$ 

```

The interesting cases of **sound** are the ones for **hyp** and $_ \rightarrow _$. Hypotheses can simply be looked up in the list of valid assumptions (using **lookup** : $\text{All } P xs \rightarrow x \in xs \rightarrow P x$), and function spaces can be handled by defining a cyclic proof:

```

sound valid (hyp  $h$ ) = lookup valid  $h$ 
sound valid  $(\tau_1 \leq \sigma_1 \rightarrow \sigma_2 \leq \tau_2) = \text{proof}$ 
where proof =  $\#$  sound (proof :: valid)  $\tau_1 \leq \sigma_1 \rightarrow$ 
              $\#$  sound (proof :: valid)  $\sigma_2 \leq \tau_2$ 

```

Note that the corecursive calls extend the list of valid assumptions with the proof currently being defined. The definition of **proof** is not guarded, but it would be if **sound** were a constructor. Yet again this problem can be solved by using the technique from Section 3. The program and WHNF types can be defined mutually as follows:

```

mutual
data  $\_ \leq_P \_ : \text{Ty } n \rightarrow \text{Ty } n \rightarrow \text{Set}$  where
  sound : All (Valid  $\_ \leq_W \_$ )  $A \rightarrow A \vdash \sigma \leq \tau \rightarrow \sigma \leq_P \tau$ 
data  $\_ \leq_W \_ : \text{Ty } n \rightarrow \text{Ty } n \rightarrow \text{Set}$  where
  done :  $\sigma \leq \tau \rightarrow \sigma \leq_W \tau$ 
   $\_ \rightarrow \_$  :  $\infty (\tau_1 \leq_P \sigma_1) \rightarrow \infty (\sigma_2 \leq_P \tau_2) \rightarrow$ 
              $\sigma_1 \rightarrow \sigma_2 \leq_W \tau_1 \rightarrow \tau_2$ 
  trans :  $\tau_1 \leq_W \tau_2 \rightarrow \tau_2 \leq_W \tau_3 \rightarrow \tau_1 \leq_W \tau_3$ 

```

The cases of **sound** listed above are now part of the **whnf** function:

```

whnf :  $\sigma \leq_p \tau \rightarrow \sigma \leq_W \tau$ 
...
whnf (sound valid (hyp h)) = lookup valid h
whnf (sound valid ( $\tau_1 \leq \sigma_1 \rightarrow \sigma_2 \leq \tau_2$ )) = proof
  where proof =  $\#$  sound (proof :: valid)  $\tau_1 \leq \sigma_1 \rightarrow$ 
                $\#$  sound (proof :: valid)  $\sigma_2 \leq \tau_2$ 

```

Note that *proof* is now guarded.

We have not found a proof for completeness of $_ \vdash _ \leq _$ with respect to $_ \leq _$ which does not use a decision procedure for subtyping. This is not entirely surprising: such a completeness proof must turn a potentially infinite proof of $\sigma \leq \tau$ into a finite proof of $[\] \vdash \sigma \leq \tau$, so some “trick” is necessary. With a suitably formulated decision procedure at hand the trick is simple: the decision procedure gives either a proof of $[\] \vdash \sigma \leq \tau$, or a proof which shows that $\sigma \leq \tau$ is impossible. In the first case we are done, and in the second case a contradiction can be derived.

4.5 Postulating an admissible rule may not be sound

Given an inductively defined inference system one can add a new rule corresponding to an admissible property without changing the set of derivable properties. It is easy to prove this statement by defining functions which translate between the two inference systems. Translating derivations from the smaller to the larger inference system is trivial. When translating in the other direction one can replace all occurrences of the new rule with instances of the proof of admissibility; this process can be implemented using recursion over the structure of the input derivation.

However, when coinduction comes into the picture this property is no longer true (de Vries 2009). The proof given above breaks down because there is no guarantee that the second translation can be implemented in a productive way. The problem is that, although the admissible rule has a proof, this proof may not be sufficiently “contractive” (for instance, the proof may replace coinductive rules in the input derivation with inductive rules in the output derivation).

The following example illustrates the problem. Capretta (2005) defines the so-called partiality monad as follows:

```

data  $\_ \nu$  (A : Set) : Set where
  return : A  $\rightarrow$   $A \nu$ 
  step  $\nu$  :  $\infty$  (A  $\nu$ )  $\rightarrow$   $A \nu$ 

```

The constructor *return* returns a result, and *step* postpones a computation. Two computations are deemed equivalent if they yield the same result, perhaps with finite differences in delay:

```

data  $\_ \approx \_$  : A  $\nu$   $\rightarrow$  A  $\nu$   $\rightarrow$  Set where
  return : return  $\nu \approx$  return  $\nu$ 
  step  $\nu$  :  $\infty$  ( $\nu$  x  $\approx$   $\nu$  y)  $\rightarrow$  step x  $\approx$  step y
  stepr : x  $\approx$   $\nu$  y  $\rightarrow$  x  $\approx$  step y
  stepl :  $\nu$  x  $\approx$  y  $\rightarrow$  step x  $\approx$  y

```

It can be proved that this equality is an equivalence relation, and that it is not trivial (assuming that *A* is inhabited). Let us now add transitivity as an inductive rule:

```

...
trans : x  $\approx$  y  $\rightarrow$  y  $\approx$  z  $\rightarrow$  x  $\approx$  z

```

Given this new constructor we can prove, using guarded coinduction, that the relation is trivial:

```

trivial : (x y : A  $\nu$ )  $\rightarrow$  x  $\approx$  y
trivial x y =
  trans (stepr (refl x))
        (trans (step ( $\#$  trivial x y)) (stepl (refl y)))

```

(Here *refl* is a proof of reflexivity.)

This problem does not affect the definition of subtyping given above, which has been proved to be equivalent to other definitions from the literature. However, it means that one should exercise caution when defining relations using mixed induction and coinduction, and avoid relying on results or intuitions which are only valid in the inductive case.

5. Total parser combinators

This section describes how parser combinators which guarantee termination can be implemented. For simplicity only recognition is discussed. The accompanying code uses the ideas presented below to implement monadic parser combinators which return proper results, not just yes or no.

5.1 Parsers

The aim is to define a data type with the following basic combinators as constructors: \emptyset , which always fails; ε , which accepts the empty string; *tok*, which accepts a given token; $_ \perp _$, symmetric choice; and $_ \cdot _$, sequencing. Another combinator is added below.

Let us first consider whether the combinators should be read inductively or coinductively. An infinite choice cannot be decided, so we choose to read choices inductively. The situation is a bit trickier for sequencing. Consider definitions like $p = p \cdot p'$ or $p = p' \cdot p$. If p' accepts the empty string, then it seems hard to make any progress with these definitions. However, if p' is guaranteed not to accept the empty string then parsing can be implemented by a structurally recursive top-down approach, because the (finite) input string will be shorter when we descend on the recursive occurrence of p . To make use of this idea we will indicate whether or not a parser is nullable (accepts the empty string) in its type, and the left (right) argument of $_ \cdot _$ will be coinductive iff the right (left) argument is not nullable. This “conditional coinduction” is encoded using the following data type:

```

data  $\infty^?$  (A : Set) : Bool  $\rightarrow$  Set where
  <_> : A  $\rightarrow$   $\infty^?$  A true
  <<_>> :  $\infty$  A  $\rightarrow$   $\infty^?$  A false

```

For convenience the index *true* is used for the inductive case, and *false* for the coinductive case. The type comes with corresponding conditional delay and force operators:

```

 $\#^?$  :  $\forall$  {b A}  $\rightarrow$  A  $\rightarrow$   $\infty^?$  A b
 $\#^?$  {true} x = < x >
 $\#^?$  {false} x = <<  $\#$  x >>
 $\flat^?$  :  $\forall$  {b A}  $\rightarrow$   $\infty^?$  A b  $\rightarrow$  A
 $\flat^?$  < x > = x
 $\flat^?$  << x >> =  $\flat$  x

```

Based on the observations and definitions above the type *P* of parsers (recognisers) can now be defined:

```

data P : Bool  $\rightarrow$  Set where
   $\emptyset$  : P false
   $\varepsilon$  : P true
  tok : Tok  $\rightarrow$  P false
   $\_ \perp \_$  :  $\forall$  {n1 n2}  $\rightarrow$  P n1  $\rightarrow$  P n2  $\rightarrow$  P (n1  $\vee$  n2)
   $\_ \cdot \_$  :  $\forall$  {n1 n2}  $\rightarrow$ 
     $\infty^?$  (P n1) n2  $\rightarrow$   $\infty^?$  (P n2) n1  $\rightarrow$  P (n1  $\wedge$  n2)

```

(*Tok* is the token type.) Note how the conditional coinduction operator is used to express whether the two arguments to the sequencing operator should be read inductively or coinductively. Note also that \emptyset and *tok* do not accept the empty string, while ε does. A choice $p_1 \mid p_2$ is nullable if either p_1 or p_2 is, and a sequence $p_1 \cdot p_2$ is nullable if both p_1 and p_2 are.

5.2 Examples

Using the definition above it is easy to define parsers which are both left and right recursive:

```
leftRight : P false
leftRight = « # leftRight » · « # leftRight »
```

Using the semantics in Section 5.3 it is easy to show that *leftRight* does not accept any string (so \emptyset could be a derived combinator).

As a more useful example of how the combinators above can be used to define derived parsers, consider the following definition of the Kleene star:

```
mutual
  _* : P false → P true
  p * = ε | p +
  _+ : P false → P false
  p + = ⟨ p ⟩ · « # (p *) »
```

The recogniser $p *$ accepts zero or more occurrences of whatever p accepts, and $p +$ accepts one or more occurrences; this is easy to prove using the semantics in Section 5.3. Note that this definition is guarded, and hence productive. Note also that p must not accept the empty string, because if it did, then the right hand side of $p +$ would have to be written $\langle p \rangle \cdot \langle p * \rangle$, which would make the definition unproductive.

One might argue that the types are a bit restrictive here: why should the argument parser not be allowed to be nullable? In the case of recognition this would make sense, but when parsers return results the situation is less clear. In that case a natural generalisation of the parser $p *$ returns a list of all the results of the parser p , and if p is nullable this list could be infinitely long (and highly ambiguous). One solution to this problem is to use a definition of $_*$ which returns a list of all results corresponding to *nonempty* matches. However, this behaviour might surprise a user of the parser combinator library who thought that $_*$ returned *all* results. We think it is better to provide the user with a combinator *nonempty* such that *nonempty* p is a variant of p which does not accept the empty string:

```
nonempty : ∀ {n} → P n → P false
```

The user can then make an explicit choice to throw away certain results: *(nonempty* p) $*$.

The example above is very small; larger examples can also be constructed. For instance, Danielsson and Norell (2009) construct *mixfix* operator grammars using a parser combinator library which is based on the ideas described here.

5.3 Semantics

The semantics of the parsers is defined inductively. The type $s \in p$ is inhabited iff the token string s is a member of the language defined by p :

```
data _∈_ : List Tok → P n → Set where
```

No string is a member of the language defined by \emptyset , so there is no constructor for it in $_∈_$. The empty string is recognised by ε :

```
ε : [] ∈ ε
```

The singleton t is recognised by *tok* t :

```
tok : [t] ∈ tok t
```

If s is recognised by p_1 , then it is also recognised by $p_1 | p_2$, and similarly for p_2 :

```
|l : s ∈ p1 → s ∈ p1 | p2
|r : s ∈ p2 → s ∈ p1 | p2
```

If s_1 is recognised by p_1 (suitably forced), and s_2 is recognised by p_2 (suitably forced), then the concatenation of s_1 and s_2 is recognised by $p_1 \cdot p_2$:

```
· : s1 ∈ p1 → s2 ∈ p2 → s1 ++ s2 ∈ p1 · p2
```

Finally, if a nonempty string is recognised by p , then it is also recognised by *nonempty* p :

```
nonempty : t :: s ∈ p → t :: s ∈ nonempty p
```

It is easy to show that the semantics and the nullability index agree: if $p : P n$, then $[] \in p$ iff n is equal to *true* (both directions can be proved by structural induction). Given this result it is easy to decide whether or not $[] \in p$; it suffices to inspect the index:

```
nullable? : ∀ {n} (p : P n) → Dec ([] ∈ p)
```

The type *Dec* P states that P is decidable; an element of *Dec* P is either a proof of P or a proof showing that P is impossible:

```
data Dec (P : Set) : Set where
  yes : P → Dec P
  no : ¬ P → Dec P
```

Here logical negation is represented as a function into the empty type \perp :

```
¬_ : Set → Set
¬ P = P → ⊥
```

5.4 Backend

Let us finally consider how the relation $_∈_$ can be decided, i.e. how a parser backend can be implemented for P . No attempt is made to make this backend efficient, the focus is on correctness.

The parser backend will be implemented using so-called derivatives (Brzozowski 1964). The derivative $\partial p t$ of p with respect to t is the “remainder” of p after p has matched the token t ; it should satisfy the equivalence

```
s ∈ ∂ p t ↔ t :: s ∈ p.
```

By applying the derivative operator ∂ to p and t_1 , then to $\partial p t_1$ and t_2 , and so on for every element in the input string s , one can decide if $s \in p$ is inhabited.

The new parser constructed by ∂ may not have the same nullability index as the original parser, so ∂ has the following type signature:

```
∂ : ∀ {n} (p : P n) (t : Tok) → P (∂n p t)
```

The extensional behaviour of $\partial^n : \forall \{n\} \rightarrow P n \rightarrow Tok \rightarrow Bool$ is uniquely constrained by the definition of ∂ , so the definition of ∂^n is omitted here.

The main derivative operator is implemented as follows. The combinators \emptyset and ε never accept any token, so they both have the derivative \emptyset :

```
∂ ∅ t = ∅
∂ ε t = ∅
```

The combinator *tok* t' has a non-zero derivative with respect to t iff t' and t are equal (*Tok* is assumed to come with a decision procedure $_≐?$ for definitional equality):

```
∂ (tok t') t with t' ≐? t
∂ (tok t') t | yes t' ≐ t = ε
∂ (tok t') t | no t' ≠ t = ∅
```

The derivative of *nonempty* p is the derivative of p :

```
∂ (nonempty p) t = ∂ p t
```

The derivative of a choice is the choice of the derivatives of its arguments:

$$\partial (p_1 \mid p_2) t = \partial p_1 t \mid \partial p_2 t$$

The final and most interesting case is sequencing. If p_1 is nullable, then the result is implemented as a choice, because the remainder of $p_1 \cdot \langle p_2 \rangle$ could be either the remainder of p_1 followed by p_2 , or the remainder of p_2 :

$$\begin{aligned} \partial (\langle p_1 \rangle \cdot \langle p_2 \rangle) t &= \langle \partial \langle p_1 \rangle t \rangle \cdot \#^? p_2 \mid \partial p_2 t \\ \partial (\langle \langle p_1 \rangle \rangle \cdot \langle p_2 \rangle) t &= \langle \# (\partial \langle p_1 \rangle t) \rangle \cdot \#^? p_2 \mid \partial p_2 t \end{aligned}$$

(Note that we may need to delay p_2 , depending on the nullability index of the derivative of p_1 .) If p_1 is not nullable, then the second choice above should not be included, because the first token which is accepted (if any) has to be accepted by p_1 :

$$\begin{aligned} \partial (\langle p_1 \rangle \cdot \langle \langle p_2 \rangle \rangle) t &= \langle \partial \langle p_1 \rangle t \rangle \cdot \#^? \langle p_2 \rangle \\ \partial (\langle \langle p_1 \rangle \rangle \cdot \langle \langle p_2 \rangle \rangle) t &= \langle \# (\partial \langle p_1 \rangle t) \rangle \cdot \#^? \langle p_2 \rangle \end{aligned}$$

The derivative operator ∂ is implemented using a lexicographic combination of guarded corecursion and structural recursion, and the index function ∂^n uses recursion over the *inductive* structure of the parser. Note that in the last two sequencing cases p_2 is delayed, but ∂ is not applied recursively to p_2 because p_1 is known not to accept the empty string.

It is straightforward to show that the derivative operator ∂ satisfies both directions of its specification:

$$\begin{aligned} \partial\text{-sound} &: s \in \partial p t \rightarrow t :: s \in p \\ \partial\text{-complete} &: t :: s \in p \rightarrow s \in \partial p t \end{aligned}$$

These statements can be proved by induction on the structure of the semantics.

Once the derivative operator is defined and proved correct it is easy to decide $_ \in _$:

$$\begin{aligned} _ \in? _ &: \forall \{n\} (s : List Tok) (p : P n) \rightarrow Dec (s \in p) \\ [] &\in? p = nullable? p \\ t :: s \in? p &\text{ with } s \in? \partial p t \\ t :: s \in? p &\mid \text{yes } s \in \partial p t = \text{yes } (\partial\text{-sound } s \in \partial p t) \\ t :: s \in? p &\mid \text{no } s \notin \partial p t = \text{no } (s \notin \partial p t \circ \partial\text{-complete}) \end{aligned}$$

In the case of the empty string the nullability index tells whether the string should be accepted or not, and otherwise $_ \in? _$ is recursively applied to the derivative and the tail of the string; the specification of ∂ ensures that this is correct.

6. Big-step semantics for both converging and diverging computations

This section shows how a big-step semantics which handles both converging and diverging computations can be defined without duplication of rules.

The untyped λ -calculus with an infinite set of constants can be represented as follows:

$$\begin{aligned} \text{data } Tm (n : \mathbb{N}) : Set \text{ where} \\ \text{con} &: \mathbb{N} \rightarrow Tm n \text{ -- Constant.} \\ \text{var} &: Fin n \rightarrow Tm n \text{ -- Variable.} \\ \lambda &: Tm (1 + n) \rightarrow Tm n \text{ -- Abstraction.} \\ _ \cdot _ &: Tm n \rightarrow Tm n \rightarrow Tm n \text{ -- Application.} \end{aligned}$$

Here de Bruijn indices are used to represent variables, just as in Section 4.1, and capture-avoiding substitution can be defined here as well:

$$_ [-] : \forall \{n\} \rightarrow Tm (1 + n) \rightarrow Tm n \rightarrow Tm n$$

Leroy and Grall (2009) define a call-by-value semantics for (a slight variant of) this language by using two big-step relations:

$t \Downarrow v$ means that t converges to the value v , and $t \Uparrow$ means that t diverges. The only values are constants and abstractions, and values can be turned into terms by using $\ulcorner _ \urcorner$:

$$\begin{aligned} \text{data } Value (n : \mathbb{N}) : Set \text{ where} \\ \text{con} &: \mathbb{N} \rightarrow Value n \\ \lambda &: Tm (1 + n) \rightarrow Value n \\ \ulcorner _ \urcorner &: \forall \{n\} \rightarrow Value n \rightarrow Tm n \\ \ulcorner \text{con } i \urcorner &= \text{con } i \\ \ulcorner \lambda t \urcorner &= \lambda t \end{aligned}$$

Leroy and Grall define the big-step relations roughly as follows:

$$\begin{aligned} \text{data } _ \Downarrow _ &: Tm n \rightarrow Value n \rightarrow Set \text{ where} \\ \text{val} &: \ulcorner v \urcorner \Downarrow v \\ \text{app} &: t_1 \Downarrow \lambda t \rightarrow t_2 \Downarrow v \rightarrow t [\ulcorner v \urcorner] \Downarrow v' \rightarrow \\ &\quad t_1 \cdot t_2 \Downarrow v' \\ \text{data } _ \Uparrow &: Tm n \rightarrow Set \text{ where} \\ \infty \cdot &: \infty (t_1 \Uparrow) \rightarrow t_1 \cdot t_2 \Uparrow \\ \cdot \infty &: t_1 \Downarrow v \rightarrow \infty (t_2 \Uparrow) \rightarrow t_1 \cdot t_2 \Uparrow \\ \text{app} &: t_1 \Downarrow \lambda t \rightarrow t_2 \Downarrow v \rightarrow \infty (t [\ulcorner v \urcorner] \Uparrow) \rightarrow \\ &\quad t_1 \cdot t_2 \Uparrow \end{aligned}$$

Note that convergence is defined inductively, while divergence is defined coinductively.

Leroy and Grall observe that the big-step definition of the semantics above contains some duplication, and note that this can be problematic when realistic languages with many rules are formalised. As an attempt to avoid this duplication they investigate a so-called coevaluation relation $_ \overset{\text{co}}{\Rightarrow} _$, defined coinductively roughly as follows:

$$\begin{aligned} \text{data } _ \overset{\text{co}}{\Rightarrow} _ &: Tm n \rightarrow Value n \rightarrow Set \text{ where} \\ \text{val} &: \ulcorner v \urcorner \overset{\text{co}}{\Rightarrow} v \\ \text{app} &: \infty (t_1 \overset{\text{co}}{\Rightarrow} \lambda t) \rightarrow \infty (t_2 \overset{\text{co}}{\Rightarrow} v) \rightarrow \\ &\quad \infty (t [\ulcorner v \urcorner] \overset{\text{co}}{\Rightarrow} v') \rightarrow t_1 \cdot t_2 \overset{\text{co}}{\Rightarrow} v' \end{aligned}$$

However, this relation turns out to be different from the ones defined above: there are terms which diverge according to $_ \Uparrow$, but which do not coevaluate to any value (consider the application of a diverging term to a stuck term like $\text{con } 0 \cdot \text{con } 1$).

In order to avoid the most glaring duplication—the two variants of the rule `app`—one can instead use mixed induction and coinduction. Define a semantic domain Sem as follows:

$$\begin{aligned} \text{data } Sem (n : \mathbb{N}) : Set \text{ where} \\ \perp &: Sem n \\ \text{val} &: Value n \rightarrow Sem n \end{aligned}$$

Here \perp stands for divergence, and `val v` for convergence with the value v . A semantics with only one occurrence of the rule `app` can then be defined:

$$\begin{aligned} \text{mutual} \\ \text{data } _ \overset{\infty^?}{\Rightarrow} _ &: Tm n \rightarrow Sem n \rightarrow Set \text{ where} \\ \text{val} &: \ulcorner v \urcorner \overset{\infty^?}{\Rightarrow} \text{val } v \\ \text{app} &: t_1 \overset{\infty^?}{\Rightarrow} \text{val } (\lambda t) \rightarrow t_2 \overset{\infty^?}{\Rightarrow} \text{val } v \rightarrow \\ &\quad t [\ulcorner v \urcorner] \overset{\infty^?}{\Rightarrow} s \rightarrow t_1 \cdot t_2 \overset{\infty^?}{\Rightarrow} s \\ \infty \cdot &: t_1 \overset{\infty^?}{\Rightarrow} \perp \rightarrow t_1 \cdot t_2 \overset{\infty^?}{\Rightarrow} \perp \\ \cdot \infty &: t_1 \overset{\infty^?}{\Rightarrow} \text{val } v \rightarrow t_2 \overset{\infty^?}{\Rightarrow} \perp \rightarrow t_1 \cdot t_2 \overset{\infty^?}{\Rightarrow} \perp \\ _ \overset{\infty^?}{\Rightarrow} _ &: Tm n \rightarrow Sem n \rightarrow Set \\ t \overset{\infty^?}{\Rightarrow} \perp &= \infty (t \overset{\infty^?}{\Rightarrow} \perp) \\ t \overset{\infty^?}{\Rightarrow} \text{val } v &= t \overset{\infty^?}{\Rightarrow} \text{val } v \end{aligned}$$

The function $_ \overset{\infty?}{\Rightarrow} _$ is used to ensure that coinduction is only used for diverging computations.

The semantics $_ \Rightarrow _$ has more rules than $_ \overset{\infty?}{\Rightarrow} _$, but at least it avoids the duplication of `app`, and it is correct; it is easy to show that $_ \Rightarrow _$ is equivalent to the combination of $_ \Downarrow _$ and $_ \Uparrow _$:

$$\begin{aligned} \text{sound-}\Downarrow & : t \Rightarrow \text{val } v \rightarrow t \Downarrow v \\ \text{complete-}\Downarrow & : t \Downarrow v \rightarrow t \Rightarrow \text{val } v \\ \text{sound-}\Uparrow & : t \Rightarrow \perp \rightarrow t \Uparrow \\ \text{complete-}\Uparrow & : t \Uparrow \rightarrow t \Rightarrow \perp \end{aligned}$$

The first two statements can be proved using structural induction, while the latter two statements can be proved using the previous two proofs and guarded coinduction.

The soundness and completeness proofs point to a problem with the definition of $_ \Rightarrow _$: the `app` rule comes with different proof principles depending on whether s is \perp or not, and this can bring the duplication of `app` back in proofs. However, in the context of mechanised formalisations the definition of a semantics can be much more important than proofs about the semantics: an incorrect proof will (hopefully) not be accepted by the proof checker, whereas an incorrect semantics can make the entire development pointless. Avoiding code duplication and corresponding issues with overview and maintenance can hence be more important in a semantics than in proofs related to the semantics.

The problem mentioned in the previous paragraph becomes more important if the actual implementation of functions defined over the type matters. (In the case of proofs the mere existence of a function with the right type is often enough.) As an example, consider the type $\text{List}^{\infty?} r A$ of lists which are finite when r is μ , and potentially infinite when r is ν :

data *Rec* : *Set* **where**

μ : *Rec*
 ν : *Rec*

$\overset{\infty?}{\Rightarrow}$: *Rec* \rightarrow *Set* \rightarrow *Set*

$\overset{\infty?}{\Rightarrow} \mu A = A$

$\overset{\infty?}{\Rightarrow} \nu A = \infty A$

data $\text{List}^{\infty?} r A$ (*r* : *Rec*) (*A* : *Set*) : *Set* **where**

$[]$: $\text{List}^{\infty?} r A$

$_ :: _$: $A \rightarrow \overset{\infty?}{\Rightarrow} r (\text{List}^{\infty?} r A) \rightarrow \text{List}^{\infty?} r A$

There is only one `cons` constructor, but the following definition of `map` has two cases for `cons`, because in one case structural recursion is used, and in the other case guarded corecursion:

$$\begin{aligned} \text{map} & : \forall \{r A B\} \rightarrow (A \rightarrow B) \rightarrow \text{List}^{\infty?} r A \rightarrow \text{List}^{\infty?} r B \\ \text{map } f [] & = [] \\ \text{map } \{\mu\} f (x :: xs) & = f x :: \text{map } f xs \\ \text{map } \{\nu\} f (x :: xs) & = f x :: \overset{\infty?}{\Rightarrow} \text{map } f (\overset{\infty?}{\Rightarrow} xs) \end{aligned}$$

(Note that it is possible to pattern match on implicit arguments by enclosing the patterns in braces.) It would be unfortunate if two equivalent finite lists with different *Rec* parameters were handled differently by `map`, so the definition above should be “*Rec*-parametric”:

$$\begin{aligned} \text{map-parametric} & : \\ \forall \{A B\} (f : A \rightarrow B) (xs : \text{List}^{\infty?} \mu A) & \rightarrow \\ \text{map } f (\text{lift } xs) \approx \text{lift } (\text{map } f xs) & \end{aligned}$$

Here $\text{lift} : \forall \{A\} \rightarrow \text{List}^{\infty?} \mu A \rightarrow \text{List}^{\infty?} \nu A$ is the obvious coercion and $_ \approx _$ is a suitably defined equality.

Note that the code duplication in `map` above could have been avoided if Agda had allowed other forms of recursion. It could be

useful to investigate some sort of language support for avoiding this duplication.

7. Related Work

7.1 Combining induction and coinduction

In the dependently typed core language $\Pi\Sigma$ Altenkirch and Oury (2008) use constructions very similar to ∞ , $\overset{\infty?}{\Rightarrow}$ and $\overset{\infty?}{\Downarrow}$, and this work influenced the current implementation of coinductive types in Agda. Another influence came from Setzer (2009), who proposes an approach to coinductive types more directly based on the category-theoretic notion of weakly final coalgebras.

Park (1980) uses nested induction and coinduction to define what a fair merge of two potentially infinite lists is, and notes that the fair merge is a fixpoint of a certain operator, but neither the least nor the greatest fixpoint.

Barwise (1989) discusses how one can define sets and proper classes by combining greatest and least fixpoints, and mentions that such definitions are used in situation theory.

Cousot and Cousot (1992, 2009) describe so-called *bi-inductive* definitions, which generalise inductive and coinductive definitions, and give a number of examples of their use. One of their examples is a semantics for a λ -calculus which captures both terminating and non-terminating behaviours in a single definition, without duplication of rules; the definition in Section 6 was inspired by this example. Adapting the example slightly to the development in Section 6, an operator F on $\wp(\text{Trm } n \times \text{Sem } n)$ is first defined by the following inference rules (note the similarity to the rules in Section 6):

$$\begin{aligned} \Gamma \nu \Uparrow \Rightarrow \text{val } v & \quad \frac{t_1 \Rightarrow \perp}{t_1 \cdot t_2 \Rightarrow \perp} \quad \frac{t_1 \Rightarrow \text{val } v \quad t_2 \Rightarrow \perp}{t_1 \cdot t_2 \Rightarrow \perp} \\ \frac{t_1 \Rightarrow \text{val } (\lambda t) \quad t_2 \Rightarrow \text{val } v \quad t [\Gamma \nu \Uparrow] \Rightarrow s}{t_1 \cdot t_2 \Rightarrow s} & \end{aligned}$$

These rules should neither be read inductively nor coinductively. The semantics is instead obtained as the least fixpoint of F with respect to the order $_ \sqsubseteq _$ defined by

$$X \sqsubseteq Y = X^+ \sqsubseteq Y^+ \quad \wedge \quad X^- \supseteq Y^-,$$

where

$$Z^+ = \{ (t, s) \in Z \mid s \neq \perp \}$$

$$Z^- = \{ (t, s) \in Z \mid s = \perp \}.$$

F is not monotone with respect to $_ \sqsubseteq _$ (which forms a complete lattice), so Cousot and Cousot give an explicit proof of the existence of a least fixpoint.

Process calculi supporting recursive definitions often include some notion of guardedness (Hoare 1985; Milner 1990; Prasad 1995). Giménez (1996) represents processes by their potentially infinite unfoldings, and uses mixed induction and coinduction to represent guardedness. Abstracting from the concrete example given by Giménez the idea is to read the arguments of the guarding constructs coinductively, and the arguments of the other constructs inductively.

HOLCF (Müller et al. 1999) supports marking individual constructor arguments as lazy. This corresponds closely to our use of ∞ , but in a domain-theoretic setting with partial functions.

Levy (2006) uses nested induction and coinduction to generalise Howe’s method to languages with non-well-founded syntax.

As mentioned in Section 2 Hancock et al. (2009) represent stream processors, i.e. functions from streams to streams, using nested induction and coinduction. Barthe et al. (2004) and Abel (2009) discuss total λ -calculi with nested least and greatest fix-

points and sized types, and Abel also shows how stream processors can be implemented in his calculus.

All uses of mixed induction and coinduction in previous sections can be phrased using an outer greatest fixpoint and an inner least fixpoint. Raffalli (1994) also discusses the other situation, with an outer least fixpoint and an inner greatest fixpoint. He defines two types of infinite bit streams, essentially

$$\begin{aligned} \mu O. \nu Z. 0 : Z + 1 : O \text{ and} \\ \nu Z. \mu O. 0 : Z + 1 : O \end{aligned}$$

(read $0 : Z + 1 : O$ as a labelled sum with constructors 0 and 1). The first one contains streams with an infinite number of zeros and a finite number of ones, while the second one contains streams with an infinite number of zeros and either a finite or an infinite number of ones.

More complicated nesting of fixpoints is also possible. Hensel and Jacobs (1997) discuss proof principles for nested induction and coinduction with arbitrary alternation of fixpoints, and include four examples: trees with finite or potentially infinite depth and finite or potentially infinite branching. Bradfield and Stirling (2007) discuss alternating fixpoints in the context of modal μ -calculus.

As an aside it is interesting to note that Agda cannot, in general, directly represent fixpoints which do not have the form $\nu X. \mu Y. F X Y$ (or an equivalent form like $\nu X. \mu Y. \mu Z. F X Y Z$). Based on the category theoretical explanation of ∞ given in Section 2 one might believe that the following definition of O should be isomorphic to $\mu O. \nu Z. 0 : Z + 1 : O$:

```
data Z (O : Set) : Set where
  0 :  $\infty (Z O) \rightarrow Z O$ 
  1 :  $O \rightarrow Z O$ 

data O : Set where
   $\downarrow$  :  $Z O \rightarrow O$ 
```

However, the following definition of a bit stream with an infinite number of ones is accepted by Agda’s productivity checker:

```
s : O
s =  $\downarrow (0 (\# 1 s))$ 
```

This is in line with the view of ∞ as a suspension type constructor. As a consequence we get that the suspension view and the category theoretical view of ∞ do not quite match. In general we have that $\mu X. \nu Y. F X Y$ is not isomorphic to $\nu Y. \mu X. F X Y$ (in the category of sets and total functions), whereas the corresponding domain theoretic expressions $\mu X. \mu Y. F X (Y_{\perp})$ and $\mu Y. \mu X. F X (Y_{\perp})$, obtained from the suspension view, are isomorphic (here \perp stands for lifting). Note, however, that this mismatch does not affect the definitions given in previous sections, because they are not of the form exhibited by O above.

7.2 Making corecursive definitions guarded

Conor McBride (personal communication) has developed a technique for ensuring guardedness, based on the work of Hancock and Setzer (2000). The idea is to represent the right-hand sides of function definitions using a type $RHS\ g$, where g indicates whether the context is guarding or not, and to only allow corecursive calls in a guarding context.

Di Gianantonio and Miculan (2003) describe a general framework for defining values using a mixture of recursion and corecursion, based on functions which satisfy a notion of contractivity. The method seems to be quite general, and has been implemented in Coq (which is based on structural recursion and guarded corecursion).

Bertot (2005) implements a filter function for streams in Coq. An unrestricted filter function is not productive, so Bertot restricts

the function’s inputs using predicates of the form “always (eventually P)”. The always part is defined coinductively, and the eventually part inductively (mixed induction and coinduction is not necessary).

The partiality monad was defined in Section 4.5. It is easy to define bind for this monad:

$$_ \gg\! = _ : \forall \{A\} \rightarrow A^{\nu} \rightarrow (A \rightarrow B^{\nu}) \rightarrow B^{\nu}$$

Unfortunately it can be inconvenient to use this definition of bind in systems based on guarded corecursion, because $_ \gg\! = _$ is not a constructor. Megacz (2007) suggests (essentially) the following alternative definition:

```
data  $\_{}^{\nu}$  (A : Set) : Set where
  return :  $A \rightarrow A^{\nu}$ 
   $\_ \gg\! = \_$  :  $\forall \{B\} \rightarrow \infty (B^{\nu}) \rightarrow (B \rightarrow \infty (A^{\nu})) \rightarrow A^{\nu}$ 
```

One can note that this corresponds directly to the first step of the technique presented in Section 3. Megacz does not translate from the second to the first type, though.

Bertot and Komendantskaya (2009) describe a method for replacing corecursion with recursion. They map values of type $Stream\ A$ to and from the isomorphic type $\mathbb{N} \rightarrow A$, and values of this type can be defined recursively. The authors state that the method is still very limited and that, as presented, it cannot handle van de Snepscheut’s corecursive definition of the Hamming numbers (Dijkstra 1981), which can easily be handled using the method described in this paper.

Morris et al. (2006) use the technique of replacing functions with constructors to show *termination* rather than productivity (see Morris et al. (2007) for an explanation of the technique). They replace a partially applied recursive call (which is not necessarily structural, because it could later be applied to anything), nested inside another recursive call, with a constructor application. If this constructor application is later encountered it is handled using structural recursion.

The technique presented here also shares some traits with Reynolds’ *defunctionalisation* (1972). Defunctionalisation is used to translate programs written in higher-order languages to first-order languages, and it basically amounts to representing function spaces using application-specific data types, and implementing interpreters for these data types.

7.3 Total parser combinators

There does not seem to be much prior work on *formally* verified termination for parser combinators (or other general parsing frameworks). McBride and McKinna (2002) define grammars inductively, and use types to ensure that a token is consumed before a non-terminal can be encountered, thereby ruling out left recursion and non-termination. Danielsson and Norell (2008) use similar ideas. Muad’Dib (2009) uses a monad annotated with Hoare-style pre- and post-conditions (Swierstra 2009) to define total parser combinators, including a fixpoint combinator whose type rules out left recursion by requiring the input to be shorter in recursive calls.

The parser combinators defined here cannot handle unrestricted left recursion (consider definitions like $p = p$). Lickman (1995) defines a parser combinator library which has a tailor-made fixpoint combinator (based on an idea due to Philip Wadler) which can be used to handle left recursion. Frost et al. (2008) also define a parser combinator library which can handle left recursion, this time by requiring the user to annotate non-terminals manually with a memoisation combinator. Both libraries give guarantees of termination (assuming that the informal proofs/arguments are correct, and that the libraries are used in the right way), and in both cases termination relies on a representation of parsers/grammars in which recursion is explicit.

Unlike all the approaches mentioned above the parser combinator library presented in Section 5 follows the example of most other parser combinator libraries in that it does not rely on any explicit representation of recursion (like grammars or fixpoints). Termination is instead guaranteed by using types to restrict how parsers can be defined. This appears to be the first time guaranteed termination has been achieved without any explicit representation of recursion.³

7.4 Big-step semantics for converging as well as diverging computations

Cousot and Cousot (2009) define several big-step semantics for an untyped λ -calculus. Every definition describes both terminating and non-terminating computations without duplication of rules.

Nakata and Uustalu (2009) define a big-step semantics for a while language. Their definition is coinductive and trace-based, and has the property that the trace can be computed (productively) for any source term, converging or diverging.

8. Conclusions

By giving a number of examples we hope to have shown that the technique of mixing induction and coinduction is generally useful, and deserves a place in the toolbox both of people formalising languages and people who are interested in ensuring totality of the programs they write.

Acknowledgments

Nils Anders Danielsson would like to thank Peter Hancock for introducing him to the technique of mixed induction and coinduction, Ulf Norell for previous joint work on total parsing, and Graham Hutton and Henrik Nilsson for feedback which improved the presentation. Thorsten Altenkirch would like to thank Nicolas Oury for joint work on $\Pi\Sigma$ which has had an impact on the work in this paper, and Graham Hutton for earlier joint but unpublished work on coinductive reasoning. Both authors would like to thank Conor McBride, Nicolas Oury and Anton Setzer for many discussions about coinduction which have influenced this work, and some anonymous reviewers for useful feedback. Both authors gratefully acknowledge funding from EPSRC; Nils Anders Danielsson would also like to thank the Royal Swedish Academy of Sciences' funds for financial support.

References

- Andreas Abel. Mixed inductive/coinductive types and strong normalization. In *Programming Languages and Systems, 5th Asian Symposium, APLAS 2007*, volume 4807 of *LNCS*, pages 286–301, 2009.
- The Agda Team. The Agda Wiki. Available at <http://wiki.portal.chalmers.se/agda/>, 2009.
- Thorsten Altenkirch and Nicolas Oury. $\Pi\Sigma$: A core language for dependently typed programming. Draft, 2008.
- Roberto M. Amadio and Luca Cardelli. Subtyping recursive types. *ACM Transactions on Programming Languages and Systems*, 15(4):575–631, 1993.
- G. Barthe, M. J. Frade, E. Giménez, L. Pinto, and T. Uustalu. Type-based termination of recursive definitions. *Mathematical Structures in Computer Science*, 14(1):97–141, 2004.
- Jon Barwise. *The Situation in Logic*, volume 17 of *CSLI Lecture Notes*, chapter Mixed Fixed Points. Center for the Study of Language and Information, Leland Stanford Junior University, 1989.

- Yves Bertot. Filters on coinductive streams, an application to Eratosthenes' sieve. In *Typed Lambda Calculi and Applications, 7th International Conference, TLCA 2005*, volume 3461 of *LNCS*, pages 102–115, 2005.
- Yves Bertot and Ekaterina Komendantskaya. Using structural recursion for corecursion. In *Types for Proofs and Programs, International Conference, TYPES 2008*, volume 5497 of *LNCS*, pages 220–236, 2009.
- Julian Bradfield and Colin Stirling. *Handbook of Modal Logic*, chapter Modal Mu-Calculi. Elsevier, 2007.
- Michael Brandt and Fritz Henglein. Coinductive axiomatization of recursive type equality and subtyping. *Fundamenta Informaticae*, 33(4):309–338, 1998.
- Janusz A. Brzozowski. Derivatives of regular expressions. *Journal of the ACM*, 11(4):481–494, 1964.
- William H. Burge. *Recursive Programming Techniques*. Addison-Wesley, 1975.
- Venanzio Capretta. General recursion via coinductive types. *Logical Methods in Computer Science*, 1(2):1–28, 2005.
- Dario Colazzo and Giorgio Ghelli. Subtyping, recursion, and parametric polymorphism in kernel Fun. *Information and Computation*, 198:71–147, 2005.
- Thierry Coquand. Infinite objects in type theory. In *Types for Proofs and Programs, International Workshop TYPES'93*, volume 806 of *LNCS*, pages 62–78, 1994.
- Patrick Cousot and Radhia Cousot. Bi-inductive structural semantics. *Information and Computation*, 207(2):258–283, 2009.
- Patrick Cousot and Radhia Cousot. Inductive definitions, semantics and abstract interpretations. In *POPL '92, Proceedings of the 19th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 83–94, 1992.
- Nils Anders Danielsson and Ulf Norell. Structurally recursive descent parsing. Presentation (given by Danielsson) at the Dependently Typed Programming workshop, Nottingham, UK, 2008.
- Nils Anders Danielsson and Ulf Norell. Parsing mixfix operators. To appear in the proceedings of the 20th International Symposium on the Implementation and Application of Functional Languages (IFL 2008), 2009.
- Edsko de Vries. Re: [Coq-Club] Adding (inductive) transitivity to weak bisimilarity not sound? (was: Need help with coinductive proof). Message to the Coq-Club mailing list, August 2009.
- Pietro Di Gianantonio and Marino Miculan. A unifying approach to recursive and co-recursive definitions. In *Types for Proofs and Programs, International Workshop, TYPES 2002*, volume 2646 of *LNCS*, pages 148–161, 2003.
- Edsger W. Dijkstra. Hamming's exercise in SASL. EWD792 (privately circulated note), 1981.
- Jon Fairbairn. Making form follow function: An exercise in functional programming style. *Software: Practice and Experience*, 17(6):379–386, 1987.
- Richard A. Frost, Rahmatullah Hafiz, and Paul Callaghan. Parser combinators for ambiguous left-recursive grammars. In *Practical Aspects of Declarative Languages, 10th International Symposium, PADL 2008*, volume 4902 of *LNCS*, pages 167–181, 2008.
- Vladimir Gapeyev, Michael Y. Levin, and Benjamin C. Pierce. Recursive subtyping revealed. *Journal of Functional Programming*, 12(6):511–548, 2002.
- Eduardo Giménez. *Un Calcul de Constructions Infinies et son Application à la Vérification de Systèmes Communicants*. PhD thesis, Ecole Normale Supérieure de Lyon, 1996.
- Andrew D. Gordon. Bisimilarity as a theory of functional programming. *Theoretical Computer Science*, 228(1–2):5–47, 1999.
- Tatsuya Hagino. *A Categorical Programming Language*. PhD thesis, University of Edinburgh, 1987.
- Peter Hancock and Anton Setzer. Interactive programs in dependent type theory. In *Computer Science Logic, 14th International Workshop, CSL 2000*, volume 1862 of *LNCS*, pages 317–331, 2000.

³Danielsson and Norell (2009) describe a less general, more specialised version of the library which is described here. Only the interface is discussed, not the backend.

- Peter Hancock, Dirk Pattinson, and Neil Ghani. Representations of stream processors using nested fixed points. *Logical Methods in Computer Science*, 5(3:9), 2009.
- Ulrich Hensel and Bart Jacobs. Proof principles for datatypes with iterated recursion. In *Category Theory and Computer Science, 7th International Conference, CTCS '97*, volume 1290 of *LNCS*, pages 220–241, 1997.
- C. A. R. Hoare. *Communicating Sequential Processes*. Prentice Hall International, 1985.
- John Hughes and Andrew Moran. Making choices lazily. In *FPCA '95, Proceedings of the seventh international conference on Functional programming languages and computer architecture*, pages 108–119, 1995.
- Xavier Leroy and Hervé Grall. Coinductive big-step operational semantics. *Information and Computation*, 207(2):284–304, 2009.
- Michael Y. Levin and Benjamin C. Pierce. TinkerType: a language for playing with formal systems. *Journal of Functional Programming*, 13(2):295–316, 2003.
- Paul Blain Levy. Infinitary Howe’s method. In *Proceedings of the Eighth Workshop on Coalgebraic Methods in Computer Science (CMCS 2006)*, volume 164 of *ENTCS*, pages 85–104, 2006.
- Paul Lickman. Parsing with fixed points. Master’s thesis, University of Cambridge, 1995.
- Conor McBride and James McKinna. Seeing and doing. Presentation (given by McBride) at the Workshop on Termination and Type Theory, Hindås, Sweden, 2002.
- Adam Megacz. A coinductive monad for Prop-bounded recursion. In *PLPV'07, Proceedings of the 2007 workshop on Programming languages meets program verification*, pages 11–20, 2007.
- Paul Francis Mendler. *Inductive Definition in Type Theory*. PhD thesis, Cornell University, 1988.
- Robin Milner. Operational and algebraic semantics of concurrent processes. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*. The MIT Press and Elsevier, 1990.
- Robin Milner and Mads Tofte. Co-induction in relational semantics. *Theoretical Computer Science*, 87(1):209–220, 1991.
- Peter Morris, Thorsten Altenkirch, and Conor McBride. Exploring the regular tree types. In *Types for Proofs and Programs, International Workshop, TYPES 2004*, volume 3839 of *LNCS*, pages 252–267, 2006.
- Peter Morris, Thorsten Altenkirch, and Neil Ghani. Constructing strictly positive families. In *Theory of Computing 2007, Proceedings of the Thirteenth Computing: The Australasian Theory Symposium (CATS2007)*, pages 111–121, 2007.
- Muad Dib. Strongly specified parser combinators. Post to the Muad Dib blog, 2009.
- Olaf Müller, Tobias Nipkow, David von Oheimb, and Oscar Slotosch. $HOLCF = HOL + LCF$. *Journal of Functional Programming*, 9(2):191–223, 1999.
- Keiko Nakata and Tarmo Uustalu. Trace-based coinductive operational semantics for While; big-step and small-step, relational and functional styles. To appear in the proceedings of TPHOLs 2009, 2009.
- Ulf Norell. *Towards a practical programming language based on dependent type theory*. PhD thesis, Chalmers University of Technology and Göteborg University, 2007.
- David Park. On the semantics of fair parallelism. In *Abstract Software Specifications*, volume 86 of *LNCS*, pages 504–526, 1980.
- K. V. S. Prasad. A calculus of broadcasting systems. *Science of Computer Programming*, 25(2–3):285–327, 1995.
- Christophe Raffalli. *L’Arithmétique Fonctionnelle du Second Ordre avec Points Fixes*. PhD thesis, Université Paris VII, 1994.
- John C. Reynolds. Definitional interpreters for higher-order programming languages. In *ACM '72, Proceedings of the ACM annual conference*, volume 2, pages 717–740, 1972.
- Anton Setzer. Coalgebras and codata in Agda. Presentation at the 3rd Wessex Theory Seminar, Bath, UK, 2009.
- Wouter Swierstra. A Hoare logic for the state monad. In *Theorem Proving in Higher Order Logics, 22nd International Conference, TPHOLs 2009*, volume 5674 of *LNCS*, pages 440–451, 2009.
- D. A. Turner. Total functional programming. *Journal of Universal Computer Science*, 10(7):751–768, 2004.
- Philip Wadler. How to replace failure by a list of successes; a method for exception handling, backtracking, and pattern matching in lazy functional languages. In *Functional Programming Languages and Computer Architecture*, volume 201 of *LNCS*, pages 113–128, 1985.
- Philip Wadler, Walid Taha, and David MacQueen. How to add laziness to a strict language, without even being odd. In *Proceedings of the 1998 ACM SIGPLAN Workshop on ML*, 1998.