Friends with Benefits

Implementing Corecursion in Foundational Proof Assistants

Jasmin Christian Blanchette[1,2], Aymeric Bouzy[3], Andreas Lochbihler[4],
Andrei Popescu[5,6], and Dmitriy Traytel[4]

[1] Vrije Universiteit Amsterdam, the Netherlands
[2] Inria Nancy – Grand Est, Nancy, France
[3] Laboratoire d'informatique, École polytechnique, Palaiseau, France
[4] Institute of Information Security, Department of Computer Science, ETH Zürich, Switzerland
[5] Department of Computer Science, Middlesex University London, UK
[6] Institute of Mathematics Simion Stoilow of the Romanian Academy, Bucharest, Romania

Abstract. We introduce AmiCo, a tool that extends a proof assistant, Isabelle/HOL, with flexible function definitions well beyond primitive corecursion. All definitions are certified by the assistant's inference kernel to guard against inconsistencies. A central notion is that of *friends*: functions that preserve the productivity of their arguments and that are allowed in corecursive call contexts. As new friends are registered, corecursion benefits by becoming more expressive. We describe this process and its implementation, from the user's specification to the synthesis of a higher-order definition to the registration of a friend. We show some substantial case studies where our approach makes a difference.

1 Introduction

Codatatypes and corecursion are emerging as a major methodology for programming with infinite objects. Unlike in traditional lazy functional programming, codatatypes support *total (co)programming* [1, 8, 30, 68], where the defined functions have a simple set-theoretic semantics and productivity is guaranteed. The proof assistants Agda [19], Coq [12], and Matita [7] have been supporting this methodology for years.

By contrast, proof assistants based on higher-order logic (HOL), such as HOL4 [64], HOL Light [32], and Isabelle/HOL [56], have traditionally provided only datatypes. Isabelle/HOL is the first of these systems to also offer codatatypes. It took two years, and about 24 000 lines of Standard ML, to move from an understanding of the mathematics [18, 67] to an implementation that automates the process of checking high-level user specifications and producing the necessary corecursion and coinduction theorems [16].

There are important differences between Isabelle/HOL and type theory systems such as Coq in the way they handle corecursion. Consider the codatatype of streams given by

```
codatatype α stream = (shd: α) ◁ (stl: α stream)
```

where ◁ (written infix) is the constructor, and shd and stl are the head and tail selectors, respectively. In Coq, a definition such as

```
corec natsFrom : nat → nat stream where
    natsFrom n = n ◁ natsFrom (n + 1)
```

which introduces the function $n \mapsto n \lhd n + 1 \lhd n + 2 \lhd \cdots$, is accepted after a syntactic check that detects the $\lhd$-guardedness of the corecursive call. In Isabelle, this check is replaced by a deeper analysis. The `primcorec` command [16] transforms a user specification into a *blueprint* object: the coalgebra $b = \lambda n.\ (n, n + 1)$. Then `natsFrom` is defined as $\mathsf{corec}_{\mathsf{stream}}\ b$, where $\mathsf{corec}_{\mathsf{stream}}$ is the fixed primitive corecursive combinator for α stream. Finally, the user specification is derived as a theorem from the definition and the characteristic equation of the corecursor.

Unlike in type theories, where (co)datatypes and (co)recursion are built-in, the HOL philosophy is to reduce every new construction to the core logic. This usually requires a lot of implementation work but guarantees that definitions introduce no inconsistencies. Since codatatypes and corecursion are derived concepts, there is no a priori restriction on the expressiveness of user specifications other than expressiveness of HOL itself.

Consider a variant of `natsFrom`, where the function add1 : nat $\rightarrow$ nat stream $\rightarrow$ nat stream adds 1 to each element of a stream:

> corec natsFrom : nat $\rightarrow$ nat stream where
>   natsFrom $n = n \lhd$ add1 (natsFrom n)

Coq's syntactic check fails on add1. After all, add1 could explore the tail of its argument before it produces a constructor, hence blocking productivity and leading to underspecification or inconsistency.

Isabelle's bookkeeping allows for more nuances. Suppose add1 has been defined as

> corec add1 : nat stream $\rightarrow$ nat stream where
>   add1 $ns = ($shd $ns + 1) \lhd$ add1 (stl ns)

When analyzing add1's specification, the `corec` command synthesizes its definition as a blueprint b. This definition can then be proved to be *friendly*, hence acceptable in corecursive call contexts when defining other functions. Functions with friendly definitions are called friendly, or *friends*. These functions preserve productivity by consuming at most one constructor when producing one.

Our previous work [17] presented the category theory underlying friends, based on more expressive blueprints than the one shown above for primitive corecursion. We now introduce a tool, AmiCo, that automates the process of applying and incrementally improving corecursion.

To demonstrate AmiCo's expressiveness and convenience, we used it to formalize eight case studies in Isabelle, featuring a variety of codatatypes and corecursion styles (Sect. 2). A few of these examples required ingenuity and suggest directions for future work. Most of the examples fall in the executable framework of Isabelle, which allows for code extraction to Haskell via Isabelle's code generator. One of them pushes the boundary of executability, integrating friends in the quantitative world of probabilities.

At the low level, the *corecursion state* summarizes what the system knows at a given point, including the set of available friends and a corecursor *up to* friends (Sect. 3). Polymorphism complicates the picture, because some friends may be available only for specific instances of a polymorphic codatatype. To each corecursor corresponds a coinduction principle up to friends and a uniqueness theorem that can be used to reason about corecursive functions. All of the constructions and theorems are derived from first

principles, without requiring new axioms or extensions of the logic. This *foundational approach* prevents the introduction of inconsistencies, such as those that have affected the termination and productivity checkers of Agda and Coq in recent years.

The user interacts with our tool via the following commands to the proof assistant (Sect. 4). The `corec` command defines a function f by extracting a blueprint *b* from a user's specification, defining f using *b* and a corecursor, and deriving the original specification from the characteristic property of the corecursor. Moreover, `corec` supports mixed recursion–corecursion specifications, exploiting proof assistant infrastructure for terminating (well-founded) recursion. Semantic proof obligations, notably termination, are either discharged automatically or presented to the user. Specifying the `friend` option to `corec` additionally registers f as a friend, enriching the corecursor state. Another command, `friend_of_corec`, registers existing functions as friendly. Friendliness amounts to the relational parametricity [60, 69] of a selected part of the definition [17], which in this paper we call a *surface*. The tool synthesizes the surface, and the parametricity proof is again either discharged automatically or presented to the user.

AmiCo is a significant piece of engineering, at about 7 000 lines of Standard ML code (Sect. 5). It subsumes a crude prototype [17] based on a shell script and template files that automated the corecursor derivation but left the blueprint and surface synthesis problems to the user. Our tool is available as part of the official Isabelle2016-1 release. The formalized examples and case studies are provided in an archive [14].

The contributions of this paper are the following:

– We describe our tool's design, algorithms, and implementation as a foundational extension of Isabelle/HOL, taking the form of the `corec`, `friend_of_corec`, `corecursive` and `coinduction_upto` commands and the *corec_unique* proof method.
– We apply our tool to a wide range of case studies, most of which are either beyond the reach of competing systems or would require type annotations and additional proofs.

More details, including thorough descriptions and proofs of correctness for the surface synthesis algorithm and the mixed recursion–corecursion pipeline, are included in a technical report [15]. Although our tool works for Isabelle, the same methodology is immediately applicable to any prover in the HOL family (including HOL4, HOL Light, HOL Zero [6], and HOL-Omega [34]), whose users represent about half of the proof assistant community. Moreover, a similar methodology is in principle applicable to provers based on type theory, such as Agda, Coq, and Matita (Sect. 6).

Conventions We recall the syntax relevant for this paper, relying on the standard set-theoretic interpretation of HOL [27].

We fix infinite sets of type variables $\alpha, \beta, \ldots$ and term variables $x, y, \ldots$ and a higher-order signature, consisting of a set of type constructors including bool and the binary constructors for functions ($\rightarrow$), products ($\times$), and sums ($+$). Types σ, τ are defined using type variables and applying type constructors, normally written postfix. Isabelle/HOL supports Haskell-style type classes, with :: expressing class membership (e.g., int :: ring).

Moreover, we assume a set of polymorphic constants c, f, g, … with declared types, including equality $= : \alpha \rightarrow \alpha \rightarrow$ bool, left and right product projections fst and snd,

and left and right sum embeddings Inl and Inr. Terms t are built from constants c and variables x by means of typed λ-abstraction and application. Polymorphic constants and terms will be freely used in contexts that require a less general type.

2 Motivating Examples

We apply AmiCo to eight case studies to demonstrate its benefits—in particular, the flexibility that friends provide and reasoning by uniqueness (of solutions to corecursive equations). The first four examples demonstrate the flexibility that friends provide. The third one also features reasoning by uniqueness. The fourth example crucially relies on a form of nested corecursion where the operator under definition must be recognized as a friend. The fifth through seventh examples mix recursion with corecursion and discuss the associated proof techniques. The last example, about a probabilistic process calculus, takes our tool to its limits: We discuss how to support corecursion through monadic sequencing and mix unbounded recursion with corecursion. All eight formalizations are available online [14], together with our earlier stream examples [17].

Since all examples are taken from the literature, we focus on the formalization with AmiCo. No detailed understanding is needed to see that they fit within the friends framework. Background information can be found in the referenced works.

Remarkably, none of the eight examples work with Coq's or Matita's standard mechanisms. Sized types in Agda [4] can cope with the first six but fail on the last two: In one case a function must inspect an infinite list unboundedly deeply, and in the other case the codatatype cannot even be defined in Agda. The Dafny verifier, which also provides codatatypes [46], supports only the seventh case study.

2.1 Coinductive Languages

Rutten [62] views formal languages as infinite tries, i.e., prefix trees branching over the alphabet with boolean labels at the nodes indicating whether the path from the root denotes a word in the language. The type α lang features corecursion through the right-hand side of the function arrow ($\rightarrow$).

$$\mathsf{codatatype}\ \alpha\ \mathsf{lang} = \mathsf{Lang}\ (o : \mathsf{bool})\ (\delta : \alpha \rightarrow \alpha\ \mathsf{lang})$$

Traytel [66] has formalized tries in Isabelle using a codatatype, defined regular operations on them as corecursive functions, and proved by coinduction that the defined operations form a Kleene algebra. Because Isabelle offered only primitive corecursion when this formalization was developed, the definition of concatenation, iteration, and shuffle product was tedious, spanning more than a hundred lines.

Corecursion up to friends eliminates this tedium. The following extract from an Isabelle formalization is all that is needed to define the main operations on languages:

$$\mathsf{corec}\ (\mathsf{friend})\ + : \alpha\ \mathsf{lang} \rightarrow \alpha\ \mathsf{lang} \rightarrow \alpha\ \mathsf{lang}\ \mathsf{where}$$
$$L + K = \mathsf{Lang}\ (o\ L \vee o\ K)\ (\lambda a.\ \delta\ L\ a + \delta\ K\ a)$$

$$\mathsf{corec}\ (\mathsf{friend})\ \cdot : \alpha\ \mathsf{lang} \rightarrow \alpha\ \mathsf{lang} \rightarrow \alpha\ \mathsf{lang}\ \mathsf{where}$$
$$L \cdot K = \mathsf{Lang}\ (o\ L \wedge o\ K)\ (\lambda a.\ \mathsf{if}\ o\ L\ \mathsf{then}\ (\delta\ L\ a \cdot K) + \delta\ K\ a\ \mathsf{else}\ \delta\ L\ a \cdot K)$$

4

```
corec (friend) * : α lang → α lang where
   L* = Lang True (λa. δ L a · L*)

corec (friend) ‖ : α lang → α lang → α lang where
   L ‖ K = Lang (o L ∧ o K) (λa. (δ L a · K) + (L · δ K a))
```

Concatenation (·) and shuffle product (‖) are corecursive up to alternation (+), and iteration (*) is corecursive up to concatenation (·). All four definitions use an alternative λ-based syntax for performing corecursion under the right-hand side of →, instead of applying the functorial action $map_{\rightarrow} = \circ$ (composition) associated with →.

The `corec` command is provided by AmiCo, whereas `codatatype` and `primcorec` (Sect. 3.2) has been part of Isabelle since 2013. The `friend` option registers the defined functions as friends and automatically discharges the emerging proof obligations, which ensure that friends consume at most one constructor to produce one constructor.

Proving equalities on tries conveniently works by coinduction up to congruence (Sect. 3.7). Already before `corec`'s existence, Traytel was able to write automatic one-line proofs such as

```
lemma K · (L + M) = K · L + K · M
   by (coinduction arbitrary: K L M rule: +.coinduct) auto
```

The *coinduction* proof method [16] instantiates the bisimulation witness of the given coinduction rule before applying it backwards. Without `corec`, the rule +.*coinduct* of coinduction up to congruence had to be stated and proved manually, including the manual inductive definition of the congruence closure under +.

Overall, the usage of `corec` compressed Traytel's development from 750 to 600 lines of Isabelle text. In Agda, Abel [3] has formalized Traytel's work up to proving the recursion equation $L^* = \varepsilon + L \cdot L^*$ for iteration (*) in 219 lines of Agda text, which correspond to 125 lines in our version. His definitions are as concise as ours, but his proofs require more manual steps.

2.2 Knuth–Morris–Pratt String Matching

Building on the trie view of formal languages, van Laarhoven [44] discovered a concise formulation of the Knuth–Morris–Pratt algorithm [41] for finding one string in another:

```
is_substring_of xs ys = match (mk_table xs) ys

match t xs = (o t ∨ (xs ≠ [] ∧ match (δ t (hd x) (tl xs))))

mk_table xs = let table = tab xs (λ_. table) in table

tab [] f        = Lang True f
tab (x ◁ xs) f = Lang False (λc. if c = x then tab xs (δ (f x)) else f c)
```

Here, we overload the stream constructor ◁ for finite lists; hd and tl are the selectors. In our context, table : α lang is the most interesting definition because it corecurses through tab. Since there is no constructor guard, table would appear not to be productive. However, the constructor is merely hidden in tab and can be pulled out by unrolling the definition of tab as follows.

As the first step, we register Δ defined by $\Delta\ xs\ f = \delta\ (\text{tab}\ xs\ f)$ as a friend, using the `friend_of_corec` command provided by our tool. The registration of an existing function as a friend requires us to supply an equation with a constructor-guarded right-hand side and to prove the equation and the parametricity of the destructor-free part of the right-hand side, called the surface (Sect. 3.4). Then the definition of table corecurses through Δ. Finally, we derive the original specification by unrolling the definition. We can use the derived specification in the proofs, because proofs in HOL do not depend on the actual definition (unlike in type theory).

```
corec tab : α list → (α → α lang) → α lang where
  tab xs f = Lang (xs = []) (λc. if xs = [] ∨ hd xs ≠ c then f c else tab (tl xs) (δ (f c)))

definition Δ : α list → (α → α lang) → α → α lang where
  Δ xs f = δ (tab xs f)

friend_of_corec Δ where
  Δ xs f c = Lang
    (if xs = [] ∨ hd xs ≠ c then o (f x) else tl xs = [])
    (if xs = [] ∨ hd xs ≠ c then δ (f x) else Δ (tl xs) (δ (f c)))
  ⟨two-line proof of the equation and of parametricity⟩

context fixes xs : α list begin
  corec table : α lang where
    table = Lang (xs = []) (Δ xs (λ_. table))

  lemma table = tab xs (λ_. table)
    ⟨one-line proof⟩
end
```

2.3   The Stern–Brocot Tree

The next application involves infinite trees of rational numbers. It is based on Hinze's work on the Stern–Brocot and Bird trees [33] and the Isabelle formalization by Gammie and Lochbihler [25]. It illustrates reasoning by uniqueness (Sect. 3.7).

The Stern–Brocot tree contains all the rational numbers in their lowest terms. It is an infinite binary tree frac tree of formal fractions frac = nat × nat. Each node is labeled with the median of its rightmost and leftmost ancestors, where median $(a, c)\ (b, d) = (a + b, c + d)$. Gammie and Lochbihler define the tree via an iterative helper function.

```
codatatype α tree = Node (root: α) (left: α tree) (right: α tree)

primcorec stern_brocot_gen : frac → frac → frac tree where
  stern_brocot_gen l u =
    let m = median l u in Node m (stern_brocot_gen l m) (stern_brocot_gen m u)

definition stern_brocot : frac tree where
  stern_brocot = stern_brocot_gen (0, 1) (1, 0)
```

Using AmiCo, we can directly formalize Hinze's corecursive specification of the tree, where nxt $(m, n) = (m + n, n)$ and swap $(m, n) = (n, m)$. The tree is corecursive up to the two friends suc and $1 / t$.

```
corec (friend) suc : frac tree → frac tree where
   suc t = Node (nxt (root t)) (suc (left t)) (suc (right t))
```

```
corec (friend) 1 / _ : frac tree → frac tree where
   1 / t = Node (swap (root t)) (1 / left t) (1 / right t)
```

```
corec stern_brocot : frac tree where
   stern_brocot = Node (1, 1) (1 / (suc (1 / stern_brocot))) (suc stern_brocot)
```

Without the iterative detour, the proofs, too, become more direct as the statements need not be generalized for the iterative helper function. For example, Hinze relies on the uniqueness principle to show that a loopless linearization stream stern_brocot of the tree yields Dijkstra's fusc function [23] given by

$$\mathsf{fusc} = 1 \lhd \mathsf{fusc}' \qquad \mathsf{fusc}' = 1 \lhd (\mathsf{fusc} + \mathsf{fusc}' - 2 \cdot (\mathsf{fusc} \bmod \mathsf{fusc}'))$$

where all arithmetic operations are lifted to streams elementwise—e.g., $xs + ys = \mathsf{map}_{\mathsf{stream}} (+) (xs \,\zipf\, ys)$, where $\zipf$ zips two streams. We define fusc and stream as follows. To avoid the mutual corecursion, we inline fusc in fusc' for the definition with `corec`, after having registered the arithmetic operations as friends:

```
corec fusc' : nat stream where
   fusc' = 1 ◁ ((1 ◁ fusc') + fusc' − 2 · ((1 ◁ fusc') mod fusc'))
```

```
definition fusc : nat stream where
   fusc = 1 ◁ fusc'
```

```
corec chop : α tree → α tree where
   chop (Node x l r) = Node (root l) r (chop l)
```

```
corec stream : α tree → α stream where
   stream t = root t ◁ stream (chop t)
```

Hinze proves that stream stern_brocot equals $\mathsf{fusc} \,\zipf\, \mathsf{fusc}'$ by showing that both satisfy the corecursion equation $x = (1, 1) \lhd \mathsf{map}_{\mathsf{stream}} \mathsf{step}\ x$, where $\mathsf{step}\ (m, n) = (n, m + n - 2 \cdot (m \bmod n))$. This equation yields the loopless algorithm, because siterate step $(1, 1)$ satisfies it as well, where siterate is defined by

```
primcorec siterate : (α → α) → α → α stream where
   siterate f x = x ◁ siterate f (f x)
```

Our tool generates a proof rule for uniqueness of solutions to the recursion equation (Sect. 3.7). We conduct the equivalence proofs using this rule.

For another example, all rational numbers also occur in the Bird tree given by

```
corec bird : frac tree where
   bird = Node (1, 1) (1 / suc bird) (suc (1 / bird))
```

It satisfies $1 / \mathsf{bird} = \mathsf{mirror}\ \mathsf{bird}$, where mirror corecursively swaps all subtrees. Again, we prove this identity by showing that both sides satisfy the corecursion equation $x = \mathsf{Node}\ (1, 1)\ (\mathsf{suc}\ (1 / x))\ (1 / \mathsf{suc}\ x)$. This equation does not correspond to any function defined with `corec`, but we can derive its uniqueness principle using our proof method *corec_unique* without defining the function. The Isabelle proof is quite concise:

```
let ?H = λx. Node (1, 1) (suc (1 / x)) (1 / suc x)
have mb: mirror bird = ?H (mirror bird)        by (rule tree.expand) …
have unique: ∀t. t = ?H t ⟶ t = mirror bird    by corec_unique (fact mb)
have 1 / bird = ?H (1 / bird)                   by (rule tree.expand) …
then show 1 / bird = mirror bird                by (rule unique)
```

No coinduction is needed: The identities are proved by expanding the definitions a finite number of times (once each here). We also show that odd_mirror bird = stern_brocot by uniqueness, where odd_mirror swaps the subtrees only at levels of odd depth.

Gammie and Lochbihler manually derive each uniqueness rule using a separate coinduction proof. For odd_mirror alone, the proof requires 25 lines. With AmiCo's *corec_unique* proof method, such proofs are automatic.

2.4 Breadth-First Tree Labeling

Abel and Pientka [4] demonstrate the expressive power of sized types in Agda with the example of labeling the nodes of an infinite binary tree in breadth-first order, which they adapted from Jones and Gibbons [39]. The function bfs takes a stream of streams of labels as input and labels the nodes at depth i according to a prefix of the ith input stream. It also outputs the streams of unused labels. Then bf ties the knot by feeding the unused labels back into bfs:

$$\text{bfs } ((x \lhd xs) \lhd ys) =$$
$$\quad \text{let } (l, ys') = \text{bfs } ys; (r, ys'') = \text{bfs } ys' \text{ in } (\text{Node } x\ l\ r, xs \lhd ys'')$$
$$\text{bf } xs = \text{let } (t, lbls) = \text{bfs } (xs \lhd lbls) \text{ in } t$$

Because bfs returns a pair, we define the two projections separately and derive the original specification for bfs trivially from the definitions. One of the corecursive calls to bfs_2 occurs in the context of bfs_2 itself—it is "self-friendly" (Sect. 4.2).

```
corec (friend) bfs₂ : α stream stream → α stream stream
  where bfs₂ ((x ◁ xs) ◁ ys) = xs ◁ bfs₂ (bfs₂ ys)

corec bfs₁ : α stream stream → α tree where
  bfs₁ ((x ◁ xs) ◁ ys) = Node x (bfs₁ ys) (bfs₁ (bfs₂ ys))

definition bfs : α stream → α tree where
  bfs xss = (bfs₁ xss, bfs₂ xss)

corec labels : α stream → α stream stream where
  labels xs = bfs₂ (xs ◁ labels xs)

definition bf : α stream → α tree where
  bf xs = bfs₁ (xs ◁ labels xs)
```

For comparison, Abel's and Pientka's formalization in Agda is of similar size, but the user must provide some size hints for the corecursive calls.

2.5 Stream Processors

Stream processors are a standard example of mixed fixpoints:

```
datatype (α,β,δ) sp_μ = Get (α → (α,β,δ) sp_μ) | Put β δ
codatatype (α,β) sp_ν = In (out: (α,β, (α,β) sp_ν) sp_μ)
```

When defining functions on these objects, we previously had to break them into a recursive and a corecursive part, using Isabelle's `primcorec` command for the latter [16]. Since our tool supports mixed recursion–corecursion, we can now express functions on stream processors more directly.

We present two functions. The first one runs a stream processor:

```
corecursive run : (α,β) sp_ν → α stream → β stream where
    run sp s = case out sp of
        Get f ⇒ run (In (f (shd s))) (stl s)
      | Put b sp ⇒ b ◁ run sp s
    ⟨two-line termination proof⟩
```

The second function, ∘∘, composes two stream processors:

```
corec (friend) get where
    get f = In (Get (λa. out (f a)))

corecursive ∘∘ : (β,γ) sp_ν → (α,β) sp_ν → (α,γ) sp_ν where
    sp ∘∘ sp′ = case (out sp, out sp′) of
        (Put b sp, _) ⇒ In (Put b (sp ∘∘ sp′))
      | (Get f, Put b sp′) ⇒ In (f b) ∘∘ sp′
      | (_, Get f′) ⇒ get (λa. sp ∘∘ In (f′ a))
    ⟨two-line termination proof⟩
```

The selector out in the noncorecursive friend get is legal, because get also adds a constructor. In both cases, the `corecursive` command emits a termination proof obligation, which we discharged in two lines, using the same techniques as when defining recursive functions. This command is equivalent to `corec`, except that it lets the user discharge proof obligations instead of applying some standard proof automation.

2.6 A Calculator

Next, we formalize a calculator example by Hur et al. [37]. The calculator inputs a number, computes the double of the sum of all inputs, and outputs the current value of the sum. When the input is 0, the calculator counts down to 0 and starts again. Hur et al. implement two versions, f and g, in a programming language embedded deeply in Coq and prove that f simulates g using parameterized coinduction.

We model the calculator in a shallow fashion as a function from the current sum to a stream processor for nats. Let calc abbreviate $nat → (nat, nat) sp_ν$. We can write the program directly as a function and very closely to its specification [37, Figure 2]. In f and g, the corecursion goes through the friends get and restart, and the constructor guard is hidden in the abbreviation put x sp = In (Put x sp).

```
corec (friend) restart : calc → calc where
   restart h n = if n > 0 then put n (restart h (n − 1)) else h 0

corec f : calc where
   f n = put n (get (λv. if v ≠ 0 then f (2 · v + n) else restart f (v + n)))

corec g : calc where
   g m = put (2 · m) (get (λv. if v = 0 then restart g (2 · m) else g (v + m)))
```

Our task is to prove that g m simulates f $(2 \cdot m)$. In fact, the two can even be proved to be bisimilar. In our shallow embedding, bisimilarity coincides with equality. We can prove g m = f $(2 \cdot m)$ by coinduction with the rule generated for the friends get and restart.

2.7 Lazy List Filtering

A classic example requiring a mix of recursion and corecursion is filtering on lazy lists. Given the polymorphic type of lazy lists

```
codatatype α llist = [] | (lhd: α) ◁ (ltl: α llist)
```

the task is to define the function lfilter : $(\alpha \rightarrow \text{bool}) \rightarrow \alpha$ llist $\rightarrow \alpha$ llist that retains only the elements that satisfy the given predicate. Paulson [58] defined lfilter using an inductive search predicate. His development culminates in a proof of

$$\text{lfilter } P \circ \text{lfilter } Q = \text{lfilter } (\lambda x. \ P \ x \wedge Q \ x) \tag{1}$$

In Dafny, Leino [45] suggests a definition that mixes recursion and corecursion. We can easily replicate Leino's definition in Isabelle, where set converts lazy lists to sets:

```
corecursive lfilter : (α → bool) → α llist → α llist where
   lfilter P xs = if ∀x ∈ set xs. ¬ P x then []
                  else if P (lhd xs) then lhd xs ◁ lfilter P (ltl xs)
                  else lfilter P (ltl xs)
   ⟨13-line termination proof⟩
```

The nonexecutability of the infinite ∀ quantifier in the 'if' condition is unproblematic in HOL, which has no built-in notion of computation. Lochbihler and Hölzl [48] define lfilter as a least fixpoint in the prefix order on llist. Using five properties, they substantiate that fixpoint induction leads to shorter proofs than Paulson's approach.

We show how to prove three of their properties using our definition, namely (1) and

$$\text{lfilter } P \ xs = [] \longleftrightarrow (\forall x \in \text{set } xs. \ \neg \ P \ xs) \tag{2}$$
$$\text{set } (\text{lfilter } P \ xs) = \text{set } xs \cap \{x \mid P \ x\} \tag{3}$$

We start with (2). We prove the interesting direction, $\longrightarrow$, by induction on $x \in$ set xs, where the inductive cases are solved automatically. For (3), the $\supseteq$ direction is also a simple induction on set. The other direction requires two nested inductions: first on $x \in$ set (lfilter P xs) and then a well-founded induction on the termination argument for the recursion in lfilter. Finally, we prove (1) using the uniqueness principle. We first derive the uniqueness rule for lfilter by a coinduction with a nested induction; this approach reflects the mixed recursive-corecursive definition of lfilter, which nests recursion inside corecursion.

```
lemma lfilter_unique:
  (∀xs. f xs = if ∀x ∈ set xs. ¬ P x then []
                else if P (lhd xs) then lhd xs ◁ f (ltl xs)
                else f (ltl xs)) ⟶
   f = lfilter P
```

(Our tool does not yet generate uniqueness rules for mixed recursive–corecursive definitions.) Then the proof of (1) is automatic:

```
lemma lfilter P ∘ lfilter Q = lfilter (λx. P x ∧ Q x)
  by (rule lfilter_unique) (auto elim: llist.set_cases)
```

Alternatively, we could have proved (1) by coinduction with a nested induction on the termination argument. The uniqueness principle works well because it incorporates both the coinduction and the induction. This underlines that uniqueness can be an elegant proof principle for mixed recursive–corecursive definitions, despite being much weaker than coinduction in the purely corecursive case. Compared with Lochbihler and Hölzl's proofs by fixpoint induction, our proofs are roughly of the same length, but `corecursive` eliminates the need for the lengthy setup for the domain theory.

2.8 Generative Probabilistic Values

Our final example relies on a codatatype that fully exploits Isabelle's modular datatype architecture built on bounded natural functors (Sect. 3.1) and that cannot be defined easily, if at all, in other systems. This example is covered in more detail in the report [15].

Lochbihler [47] proposes generative probabilistic values (GPVs) as a semantic domain for probabilistic input–output systems. Conceptually, each GPV chooses probabilistically between failing, terminating with a result of type α, and continuing by producing an output γ and transitioning into a reactive probabilistic value (RPV), which waits for a response ρ of the environment before moving to the generative successor state. Lochbihler models GPVs as a codatatype (α, γ, ρ) gpv. He also defines a monadic language on GPVs similar to a coroutine monad and an operation inline for composing GPVs with environment converters. The definition of inline poses two challenges. First, it corecurses through the monadic sequencing operation $(\ggg)_{\mathsf{gpv}} : (\beta, \gamma, \rho)$ gpv $\to (\beta \to (\alpha, \gamma, \rho)$ gpv$) \to (\alpha, \gamma, \rho)$ gpv. Due to HOL restrictions, all type variables in a friend's signature must show up in the resulting codatatype, which is not the case for $(\ggg)_{\mathsf{gpv}}$. To work around this, we define a copy gpv' of gpv with a phantom type parameter β, register $(\ggg)_{\mathsf{gpv'}}$ as a friend, and define inline in terms of its copy on gpv'. Second, inline recurses in a non-well-founded manner through the environment converter. Since our tool supports only mixing with well-founded recursion, we mimic the tool's internal behavior using a least fixpoint operator.

Initially, Lochbihler had manually derived the coinduction rule up to $\ggg_{\mathsf{gpv}}$, which our tool now generates. However, because of the copied type, our reformulation ended up roughly as complicated as the original. Moreover, we noted that coinduction up to congruence works only for equality; for user-defined predicates (e.g., typing judgments), the coinduction rule must still be derived manually. But even though this case study is not conclusive, it demonstrates the flexibility of the framework.

3 The Low Level: Corecursor States

Starting from the primitive corecursor provided by Isabelle [16], our tool derives corecursors up to larger and larger sets of friends. The corecursion state includes the set of friends $\mathcal{F}$ and the corecursor $corec_{\mathcal{F}}$. Four operations manipulate states:

- BASE gives the first nonprimitive corecursor by registering the first friends—the constructors (Sect. 3.3);
- STEP incorporates a new friend into the corecursor (Sect. 3.4);
- MERGE combines two existing sets of friends (Sect. 3.5);
- INSTANTIATE specializes the corecursor type (Sect. 3.6).

The operations BASE and STEP have already been described in detail and with many examples in our previous paper [17]. Here, we give a brief, self-contained account of them. MERGE and INSTANTIATE are new operations whose need became apparent in the course of implementation.

3.1 Bounded Natural Functors

The mathematics behind our tool assumes that the considered type constructors are both functors and relators, that they include basic functors such as identity, constant, sum, and product, and that they are closed under least and greatest fixpoints (initial algebras and final coalgebras). The tool satisfies this requirement by employing Isabelle's infrastructure for bounded natural functors (BNFs) [16, 67]. For example, the codatatype α stream is defined as the greatest solution to the fixpoint equation $\beta \cong \alpha \times \beta$, where both the right-hand side $\alpha \times \beta$ and the resulting type α stream are BNFs.

BNFs have both a functor and a relator structure. If K is a unary type constructor, we assume the existence of polymorphic constants for the functorial action, or map function, $map_K : (\alpha \to \beta) \to \alpha\,K \to \beta\,K$ and the relational action, or relator, $rel_K : (\alpha \to \beta \to bool) \to \alpha\,K \to \beta\,K \to bool$, and similarly for n-ary type constructors. For finite lists, map_{list} is the familiar map function, and given a relation r, $rel_{list}\,r$ relates two lists of the same length and with r-related elements positionwise. While the BNFs are functors on their covariant positions, the relator structure covers contravariant positions as well.

We assume that some of the polymorphic constants are known to be (relationally) *parametric* in some type variables, in the standard sense [60]. For example, if K is a ternary relator and $c : (\alpha, \beta, \gamma)\,K$, then c is parametric in β if $rel_K\,(=)\,r\,(=)\,c\,c$ holds for all $r : \beta \to \beta' \to bool$. In a slight departure from standard practice, if a term does not depend on a type variable α, we consider it parametric in α. The map function of a BNF is parametric in all its type variables. By contrast, $= : \alpha \to \alpha \to bool$ is not parametric in α.

3.2 Codatatypes and Primitive Corecursion

We fix a codatatype J. In general, J may depend on some type variables, but we leave this dependency implicit for now. While J also may have multiple, curried constructors, it is viewed at the low level as a codatatype with a single constructor $ctor_J : J\,K_{ctor} \to J$ and a destructor $dtor_J : J \to J\,K_{ctor}$:

codatatype $J = ctor_J\,(dtor_J: J\,K_{ctor})$

The mutually inverse constructor and destructor establish the isomorphism between J and $J\,K_{ctor}$. For streams, we have $\beta\,K_{ctor} = \alpha \times \beta$, $ctor\,(h, t) = h \lhd t$, and $dtor\,xs = (shd\,xs, stl\,xs)$. Low-level constructors and destructors combine several high-level constructors and destructors in one constant each. Internally, the `codatatype` command works on the low level, providing the high-level constructors as syntactic sugar [16].

In addition, the `codatatype` command derives a primitive corecursor $corec_J : (\alpha \to \alpha\,K_{ctor}) \to \alpha \to J$ characterized by the equation $corec_J\,b = ctor \circ map_{K_{ctor}}\,(corec_J\,b) \circ b$. The `primcorec` command, provided by Isabelle, reduces a primitively corecursive specification to a plain, acyclic definition expressed using this corecursor.

3.3 Corecursion up to Constructors

We call blueprints the arguments passed to corecursors. When defining a corecursive function f, a *blueprint* for f is produced, and f is defined as the corecursor applied to the blueprint. The expressiveness of a corecursor is indicated by the codomain of its blueprint argument. The blueprint passed to the primitive corecursor must return an $\alpha\,K_{ctor}$ value—e.g., a pair $(m, x) : nat \times \alpha$ for streams of natural numbers. The remaining corecursion structure is fixed: After producing m, we proceed corecursively with x. We cannot produce two numbers before proceeding corecursively—to do so, the blueprint would have to return $(m, (n, x)) : nat \times (nat \times \alpha)$.

Our first strengthening of the corecursor allows an arbitrary number of constructors before proceeding corecursively. This process takes a codatatype J and produces an initial corecursion state $\langle \mathcal{F}, \Sigma_{\mathcal{F}}, corec_{\mathcal{F}} \rangle$, where $\mathcal{F}$ is a set of known friends, $\Sigma_{\mathcal{F}}$ is a BNF that incorporates the type signatures of known friends, and $corec_{\mathcal{F}}$ is a corecursor. We omit the set-of-friends index whenever it is clear from the context. The initial state knows only one friend, $ctor$.

BASE : $J \rightsquigarrow \langle \mathcal{F}, \Sigma_{\mathcal{F}}, corec_{\mathcal{F}} \rangle$ where
$\quad\quad \mathcal{F} = \{ctor\} \quad \alpha\,\Sigma_{\mathcal{F}} = \alpha\,K_{ctor} \quad corec_{\mathcal{F}} : (\alpha \to \alpha\,\Sigma_{\mathcal{F}}^{+}) \to \alpha \to J$

Let us define the type $\alpha\,\Sigma_{\mathcal{F}}^{+}$ used for the corecursor. First, we let $\alpha\,\Sigma_{\mathcal{F}}^{*}$ be the free monad of Σ extended with J-constant leaves:

```
datatype α Σ*_F = Oper ((α Σ*_F) Σ_F) | Var α | Cst J
```

Inhabitants of $\alpha\,\Sigma_{\mathcal{F}}^{*}$ are (*formal*) *expressions* built from variable or constant leaf nodes (Var or Cst) and a syntactic representation of the constants in $\mathcal{F}$. Writing $\boxed{ctor}$ for $Oper : (\alpha\,\Sigma_{\mathcal{F}}^{*})\,K_{ctor} \to \alpha\,\Sigma_{\mathcal{F}}^{*}$, we can build expressions such as $\boxed{ctor}\,(1, Var\,(x : \alpha))$ and $\boxed{ctor}\,(2, \boxed{ctor}\,(3, Cst\,(xs : J)))$. The type $\alpha\,\Sigma_{\mathcal{F}}^{+}$, of *guarded expressions*, is similar to $\alpha\,\Sigma_{\mathcal{F}}^{*}$, except that it requires at least one $\boxed{ctor}$ guard on every path to a Var. Formally, $\alpha\,\Sigma_{\mathcal{F}}^{+}$ is defined as $((\alpha\,\Sigma_{\mathcal{F}}^{*})\,K_{ctor})\,\Sigma_{\mathcal{F}}^{*}$, so that K_{ctor} marks the guards. To simplify notation, we will pretend that $\alpha\,\Sigma_{\mathcal{F}}^{+} \subseteq \alpha\,\Sigma_{\mathcal{F}}^{*}$.

Guarded variable leaves represent corecursive calls. Constant leaves allow us to stop the corecursion with an immediate result of type J. The polymorphism of Σ^{*} is crucial. If we instantiate α to J, we can evaluate formal expressions with the function $eval : J\,\Sigma^{*} \to J$ given by $eval\,(\boxed{ctor}\,x) = ctor\,(map_{K_{ctor}}\,eval\,x)$, $eval\,(Var\,t) = t$, and $eval\,(Cst\,t) = t$. We also write $eval$ for other versions of the operator (e.g., for $J\,\Sigma^{+}$).

The corecursor's argument, the blueprint, returns guarded expressions consisting of one or more applications of $\boxed{ctor}$ before proceeding corecursively. Proceeding corecur-

sively means applying the corecursor to all variable leaves and evaluating the resulting expression. Formally:

$$\text{corec}_{\mathcal{F}}\ b = \text{eval} \circ \text{map}_{\Sigma_{\mathcal{F}}^{+}}\ (\text{corec}_{\mathcal{F}}\ b) \circ b$$

3.4 Adding New Friends

Corecursors can be strengthened to allow friendly functions to surround the context of the corecursive call. At the low level, we consider only uncurried functions.

A function $\mathsf{f} : \mathsf{J}\,\mathsf{K}_\mathsf{f} \to \mathsf{J}$ is friendly if it consumes at most one constructor before producing at least one constructor. Friendliness is captured by a mixture of two syntactic constraints and the semantic requirement of parametricity of a certain term, called the *surface*. The syntactic constraints amount to requiring that f is *expressible* using $\text{corec}_{\mathcal{F}}$, irrespective of its actual definition.

Specifically, f must be equal to $\text{corec}_{\mathcal{F}}\ b$ for some blueprint $b : \mathsf{J}\,\mathsf{K}_\mathsf{f} \to (\mathsf{J}\,\mathsf{K}_\mathsf{f})\,\Sigma^+$ that has the guarding constructor at the outermost position, and this object must be decomposable as $b = s \circ \text{map}_{\mathsf{K}_\mathsf{f}}\langle \text{id}, \text{dtor}\rangle$ for some $s : (\alpha \times \alpha\,\mathsf{K}_{\text{ctor}})\,\mathsf{K}_\mathsf{f} \to \alpha\,\Sigma^+$. The convolution operator $\langle f, g\rangle : \alpha \to \beta \times \gamma$ combines two functions $f : \alpha \to \beta$ and $g : \alpha \to \gamma$.

We call s the *surface* of b because it captures b's superficial layer while abstracting the application of the destructor. The surface s is more polymorphic than needed by the equation it has to satisfy. Moreover, s must be parametric in α. The decomposition, together with parametricity, ensures that friendly functions apply dtor at most once to their arguments and do not look any deeper—the "consumes at most one constructor" property.

STEP : $\langle \mathcal{F}, \Sigma_{\mathcal{F}}, \text{corec}_{\mathcal{F}}\rangle$ and $\mathsf{f} : \mathsf{J}\,\mathsf{K}_\mathsf{f} \to \mathsf{J}$ friendly $\rightsquigarrow \langle \mathcal{F}', \Sigma_{\mathcal{F}'}, \text{corec}_{\mathcal{F}'}\rangle$ where
   $\mathcal{F}' = \mathcal{F} \cup \{\mathsf{f}\}$ $\alpha\,\Sigma_{\mathcal{F}'} = \alpha\,\Sigma_{\mathcal{F}} + \alpha\,\mathsf{K}_\mathsf{f}$ $\text{corec}_{\mathcal{F}'} : (\alpha \to \alpha\,\Sigma_{\mathcal{F}'}^{+}) \to \alpha \to \mathsf{J}$

The return type of blueprints corresponding to $\text{corec}_{\mathcal{F}'}$ is $\Sigma_{\mathcal{F}'}^{+}$, where $\Sigma_{\mathcal{F}'}$ extends $\Sigma_{\mathcal{F}}$ with K_f. The type $\Sigma_{\mathcal{F}'}^{+}$ allows all guarded expressions of the previous corecursor but may also refer to f. The syntactic representations $\boxed{\mathsf{g}} : \alpha\,\Sigma_{\mathcal{F}}^{*}\,\mathsf{K}_\mathsf{g} \to \alpha\,\Sigma_{\mathcal{F}}^{*}$ of old friends $\mathsf{g} \in \mathcal{F}$ must be lifted to the type $(\alpha\,\Sigma_{\mathcal{F}'}^{*})\,\mathsf{K}_\mathsf{g} \to \alpha\,\Sigma_{\mathcal{F}'}^{*}$, which is straightforward. In the sequel, we will reuse the notation $\boxed{\mathsf{g}}$ for the lifted syntactic representations. In addition to $\boxed{\mathsf{g}}$, new expressions are allowed to freely use the syntactic representation $\boxed{\mathsf{f}} : (\alpha\,\Sigma_{\mathcal{F}'}^{*})\,\mathsf{K}_\mathsf{f} \to \alpha\,\Sigma_{\mathcal{F}'}^{*}$ of the new friend f, defined as $\boxed{\mathsf{f}} = \text{Oper} \circ \text{Inr}$. Like for $\boxed{\text{ctor}}$, we have $\text{eval}\,(\boxed{\mathsf{f}}\,x) = \mathsf{f}\,(\text{map}_{\mathsf{K}_\mathsf{f}}\,\text{eval}\,x)$. As before, we have $\text{corec}_{\mathcal{F}'}\ b = \text{eval} \circ \text{map}_{\Sigma_{\mathcal{F}'}^{+}}\ (\text{corec}_{\mathcal{F}'}\ b) \circ b$.

Consider the corecursive specification of pointwise addition on streams of numbers, where $\alpha\,\mathsf{K}_{\text{ctor}}$ is $\text{nat} \times \alpha$ and $\text{dtor}\ xs = (\text{shd}\ xs, \text{stl}\ xs)$:

$$xs \oplus ys = (\text{shd}\ xs + \text{shd}\ ys) \lhd (\text{stl}\ xs \oplus \text{stl}\ ys)$$

To make sense of this specification, we take $\alpha\,\mathsf{K}_\oplus$ to be $\alpha \times \alpha$ and define $\oplus$ as $\text{corec}_{\mathcal{F}}\ b$, where the blueprint b is

$$\lambda p.\ (\text{shd}\ (\text{fst}\ p) + \text{shd}\ (\text{snd}\ p)) \boxed{\lhd} \text{Var}\ (\text{stl}\ (\text{fst}\ p), \text{stl}\ (\text{snd}\ p))$$

To register $\oplus$ as friendly, we must decompose b as $s \circ \text{map}_{\mathsf{K}_\oplus}\langle \text{id}, \text{dtor}\rangle$. Expanding the definition of $\text{map}_{\mathsf{K}_\oplus}$, we get

$$\mathsf{map}_{\mathsf{K}_\oplus}\,\langle\mathsf{id},\mathsf{dtor}\rangle$$
$$= \lambda p.\,((\mathsf{fst}\,p,\mathsf{dtor}\,(\mathsf{fst}\,p)),(\mathsf{snd}\,p,\mathsf{dtor}\,(\mathsf{snd}\,p)))$$
$$= \lambda p.\,((\mathsf{fst}\,p,(\mathsf{shd}\,(\mathsf{fst}\,p),\mathsf{stl}\,(\mathsf{fst}\,p))),(\mathsf{snd}\,p,(\mathsf{shd}\,(\mathsf{snd}\,p),\mathsf{stl}\,(\mathsf{snd}\,p))))$$

It is easy to see that the following term is a suitable surface s:

$$\lambda p'.\,(\mathsf{fst}\,(\mathsf{snd}\,(\mathsf{fst}\,p'))+\mathsf{fst}\,(\mathsf{snd}\,(\mathsf{snd}\,p')))\,\boxdot\,\mathsf{Var}\,(\mathsf{snd}\,(\mathsf{snd}\,(\mathsf{fst}\,p')),\mathsf{snd}\,(\mathsf{snd}\,(\mathsf{snd}\,p')))$$

In Sect. 4, we give more details on how the system synthesizes blueprints and surfaces.

3.5 Merging Corecursion States

Most formalizations are not linear. A module may import several other modules, giving rise to a directed acyclic graph of dependencies. We can reach a situation where the codatatype has been defined in module A; its corecursor has been extended with two different sets of friends $\mathcal{F}_B$ and $\mathcal{F}_C$ in modules B and C, each importing A; and finally module D, which imports B and C, requires a corecursor that mixes friends from $\mathcal{F}_B$ and $\mathcal{F}_C$. To support this scenario, we need an operation that merges two corecursion states.

$$\mathrm{MERGE}: \quad \langle\mathcal{F}_1,\Sigma_{\mathcal{F}_1},\mathsf{corec}_{\mathcal{F}_1}\rangle \text{ and } \langle\mathcal{F}_2,\Sigma_{\mathcal{F}_2},\mathsf{corec}_{\mathcal{F}_2}\rangle \rightsquigarrow \langle\mathcal{F},\Sigma_{\mathcal{F}},\mathsf{corec}_{\mathcal{F}}\rangle \text{ where}$$
$$\mathcal{F}=\mathcal{F}_1\cup\mathcal{F}_2 \quad \alpha\,\Sigma_{\mathcal{F}}=\alpha\,\Sigma_{\mathcal{F}_1}+\alpha\,\Sigma_{\mathcal{F}_2} \quad \mathsf{corec}_{\mathcal{F}}:(\alpha\to\alpha\,\Sigma_{\mathcal{F}}^+)\to\alpha\to\mathsf{J}$$

The return type of blueprints for $\mathsf{corec}_{\mathcal{F}}$ is $\Sigma_{\mathcal{F}}^+$, where $\Sigma_{\mathcal{F}}$ is the sum of the two input signatures $\Sigma_{\mathcal{F}_1}$ and $\Sigma_{\mathcal{F}_2}$. By lifting the syntactic representations of old friends using overloading, we establish the invariant that for each $\mathsf{f}\in\mathcal{F}$ of a corecursor state, there is a syntactic representation $\boxed{\mathsf{f}}:\Sigma_{\mathcal{F}}^+\,\mathsf{K}_\mathsf{f}\to\Sigma_{\mathcal{F}}^+$. The function eval is then defined in the usual way and constitutes the main ingredient in the definition of $\mathsf{corec}_{\mathcal{F}}$ with the usual characteristic equation. For operations $\mathsf{f}\in\Sigma_{\mathcal{F}_1}\cap\Sigma_{\mathcal{F}_2}$, two syntactic representations are available; we arbitrarily choose the one inherited from $\Sigma_{\mathcal{F}_1}$.

3.6 Type Instantiation

We have so far ignored the potential polymorphism of J. Consider $\mathsf{J}=\alpha$ stream. The operations on corecursor states allow friends of type $(\alpha\text{ stream})\,\mathsf{K}\to\alpha$ stream but not $(\mathsf{nat}\text{ stream})\,\mathsf{K}\to\mathsf{nat}$ stream. To allow friends for nat stream, we must keep track of specialized corecursors. First, we need an operation for instantiating corecursor states.

$$\mathrm{INSTANTIATE}: \quad \langle F,\Sigma_{\mathcal{F}},\mathsf{corec}_{\mathcal{F}}\rangle \rightsquigarrow \langle F[\overline{\sigma}/\overline{\alpha}],\Sigma_{\mathcal{F}}[\overline{\sigma}/\overline{\alpha}],\mathsf{corec}_{\mathcal{F}}[\overline{\sigma}/\overline{\alpha}]\rangle$$

Once we have derived a specific corecursor for nat stream, we can extend it with friends of type $(\mathsf{nat}\text{ stream})\,\mathsf{K}\to\mathsf{nat}$ stream. Such friends cannot be added to the polymorphic corecursor, but the other direction works: Any friend of a polymorphic corecursor is also a friend of a specialized corecursor. Accordingly, we maintain a Pareto optimal subset of corecursor state instances $\{\langle\mathcal{F}_\mathsf{S},\Sigma_{\mathcal{F}_\mathsf{S}},\mathsf{corec}_{\mathcal{F}_\mathsf{S}}\rangle\mid\mathsf{S}\le\mathsf{J}\}$, where $\sigma'\le\sigma$ denotes that the type σ' can be obtained from the type σ by applying a type substitution.

More specific corecursors are stored only if they have more friends: For each pair of corecursor instances for S_1 and S_2 contained in the Pareto set, we have $\mathcal{F}_{\mathsf{S}_1}\supset\mathcal{F}_{\mathsf{S}_2}$ whenever $\mathsf{S}_1<\mathsf{S}_2$. All the corecursors in the Pareto set are kept up to date. If we add

a friend to a corecursor instance for S from the set via STEP, it is also propagated to all instances S′ of S by applying INSTANTIATE to the output of STEP and combining the result with the existing corecursor state for S′ via MERGE. When analyzing a user specification, corec selects the most specific applicable corecursor.

Eagerly computing the entire Pareto set is exponentially expensive. Consider a co-datatype (α, β, γ) J and the friends f for $(\mathsf{nat}, \beta, \gamma)$ J, g for $(\alpha, \beta :: \mathsf{ring}, \gamma)$ J, and h for $(\alpha, \beta, \mathsf{bool})$ J. The set would contain eight corecursors, each with a different subset of $\{\mathsf{f}, \mathsf{g}, \mathsf{h}\}$ as friends. To avoid such an explosion, we settle for a lazy derivation strategy. In the above example, the corecursor for $(\mathsf{nat}, \beta :: \mathsf{ring}, \mathsf{bool})$ J, with f, g, h as friends, is derived only if a definition needs it.

3.7 Reasoning Principles

The primary activity of a working formalizer is to develop proofs. To conveniently reason about nonprimitively corecursive functions, corec provides two reasoning principles: coinduction up to congruence and a uniqueness theorem.

Coinduction up to Congruence Codatatypes are equipped with a coinduction principle. Coinduction reduces the task of proving equality between two inhabitants l and r of a codatatype to the task of exhibiting a relation R which relates l and r and is closed under application of destructors. A relation closed under destructors is called a *bisimulation*. The codatatype command derives a plain coinduction rule. The rule for stream follows:

$$\frac{R\,l\,r \quad \forall xs\,xs'.\ R\,xs\,xs' \longrightarrow \mathsf{shd}\,xs = \mathsf{shd}\,xs' \land R\,(\mathsf{stl}\,xs)\,(\mathsf{stl}\,xs')}{l = r}$$

To reason about functions that are corecursive up to a set of friends, a principle of coinduction up to congruence of friends is crucial. For a corecursor with friends $\mathcal{F}$, our tool derives a rule that is identical to the standard rule except with $R^{\mathcal{F}}\,(\mathsf{stl}\,xs)\,(\mathsf{stl}\,xs')$ instead of $R\,(\mathsf{stl}\,xs)\,(\mathsf{stl}\,xs')$, where $R^{\mathcal{F}}$ denotes the congruence closure of the relation R with respect to the friendly operations $\mathcal{F}$.

After registering a binary $\oplus$ on nat stream as friendly, the introduction rules for the inductively defined congruence closure include

$$\frac{x = x' \quad R^{\mathcal{F}}\,xs\,xs'}{R^{\mathcal{F}}\,(x \lhd xs)\,(x' \lhd xs')} \qquad \frac{R^{\mathcal{F}}\,xs\,xs' \quad R^{\mathcal{F}}\,ys\,ys'}{R^{\mathcal{F}}\,(xs \oplus ys)\,(xs' \oplus ys')}$$

Since the tool maintains a set of incomparable corecursors, there is also a set of coinduction principles and a set of sets of introduction rules. The corec command orders the set of coinduction principles by increasing generality, which works well with Isabelle's philosophy of applying the first rule that matches.

In some circumstances, it may be necessary to reason about the union of friends associated with several incomparable corecursors. To continue with the example from Sect. 3.6, suppose we want to prove a formula about $(\mathsf{nat}, \beta :: \mathsf{ring}, \mathsf{bool})$ J by coinduction up to f, g, h before the corresponding corecursor has been derived. Users can derive it and the associated coinduction principle by invoking a dedicated command:

```
coinduction_upto (nat, β :: ring, bool) J
```

Uniqueness Principles It is sometimes possible to achieve better automation by employing a more specialized proof method than coinduction. Uniqueness principles exploit the property that the corecursor is the unique solution to a fixpoint equation:

$$h = \mathsf{eval} \circ \mathsf{map}_{\Sigma+}\, h \circ b \longrightarrow h = \mathsf{corec}_{\mathcal{F}}\, b$$

This rule can be seen as a less powerful version of coinduction, where the bisimulation relation has been preinstantiated. In category-theoretic terms, the existence and uniqueness of a solution means that we maintain on J a completely iterative algebra [51] (whose signature is gradually incremented with each additional friend).

For concrete functions defined with `corec`, uniqueness rules can be made even more precise by instantiating the blueprint b. For example, the pointwise addition on streams from Sect. 3.4

```
corec ⊕ : nat stream → nat stream → nat stream where
  xs ⊕ ys = (shd xs + shd ys) ◁ (stl xs ⊕ stl ys)
```

yields the following uniqueness principle:

$$(\forall xs\ ys.\ h\ xs\ ys = (\mathsf{shd}\ xs + \mathsf{shd}\ ys) \lhd h\ (\mathsf{stl}\ xs)\ (\mathsf{stl}\ ys)) \longrightarrow h = \oplus$$

Reasoning by uniqueness is not restricted to functions defined with `corec`. Suppose $t\,\bar{x}$ is an arbitrary term depending on a list of free variables $\bar{x}$. The *corec_unique* proof method, also provided by our tool, transforms proof obligations of the form

$$(\forall \bar{x}.\ h\,\bar{x} = H\,\bar{x}\,h) \longrightarrow h\,\bar{x} = t\,\bar{x}$$

into $\forall \bar{x}.\ t\,\bar{x} = H\,\bar{x}\,t$. The higher-order functional H must be such that the equation $h\,\bar{x} = H\,\bar{x}\,h$ would be a valid `corec` specification (but without nested calls to h or unguarded calls). Internally, *corec_unique* extracts the blueprint b from $H\,\bar{x}\,h$ as if it would define h with $\mathsf{corec}_{\mathcal{F}}$ and uses the uniqueness principle for $\mathsf{corec}_{\mathcal{F}}$ instantiated with b to achieve the described transformation.

4 The High Level: From Commands to Definitions

AmiCo's two main commands `corec` (Sect. 4.1) and `friend_of_corec` (Sect. 4.2) introduce corecursive functions and register friends. We describe synthesis algorithms for any codatatype as implemented in the tool. We also show how to capture the "consumes at most one constructor, produces at least one constructor" contract of friends.

4.1 Defining Corecursive Functions

The `corec` command reduces the user's corecursive equation to non(co)recursive primitives, so as to guard against inconsistencies. To this end, the command engages in a chain of definitions and proofs. Recall the general context:

– The codatatype J is defined as a fixpoint of a type constructor $\alpha\,\mathsf{K}_{\mathsf{ctor}}$ equipped with constructor ctor and destructor dtor.

– The current set of friends $\mathcal{F}$ contains ctor and has a signature $\Sigma_{\mathcal{F}}$ (or Σ). Each friend $f \in \mathcal{F}$ of type $J\,K_f \to J$ has a companion syntactic expression $\boxed{f} : (\alpha\,\Sigma^*)\,K_f \to \alpha\,\Sigma^*$.
– The corecursor up to $\mathcal{F}$ is $\mathrm{corec}_{\mathcal{F}} : (\alpha \to \alpha\,\Sigma^+) \to \alpha \to J$.

In general, J may be polymorphic and f may take more than one argument, but these are minor orthogonal concerns here. As before, we write $\alpha\,\Sigma^*$ for the type of formal expressions built from α-leaves and friend symbols $\boxed{f}$, and $\alpha\,\Sigma^+$ for $\boxed{\mathrm{ctor}}$-guarded formal expressions. For $\alpha = J$, we can evaluate the formal expressions into elements of J, by replacing each $\boxed{f}$ with f and omitting the Var and Cst constructors. Finally, we write eval for the evaluation functions of various types of symbolic expressions to J.

Consider the command

$$\mathtt{corec}\ g : A \to J\ \mathtt{where}\ g\,x = u_{g,x}$$

where $u_{g,x} : J$ is a term that may refer to g and x. The first task of $\mathtt{corec}$ is to synthesize a blueprint object $b : A \to A\,\Sigma^+$ such that

$$\mathrm{eval}\,(\mathrm{map}_{\Sigma^+}h\,(b\,x)) = u_{h,x} \tag{4}$$

holds for all $h : A \to J$. This equation states that the synthesized blueprint must produce, by evaluation, the concrete right-hand side of the user equation. The unknown function h represents corecursive calls, which will be instantiated to g once g is defined. To the occurrences of h in $u_{h,x}$ correspond occurrences of Var in b.

Equipped with a blueprint, we define $g = \mathrm{corec}_{\mathcal{F}}\,b$ and derive the user equation:

$$
\begin{aligned}
g\,x &= \mathrm{corec}_{\mathcal{F}}\,b\,x && \{\text{by definition of }g\} \\
&= \mathrm{eval}\,(\mathrm{map}_{\Sigma^+}(\mathrm{corec}\,b)\,(b\,x)) && \{\text{by }\mathrm{corec}_{\mathcal{F}}\text{'s equation}\} \\
&= \mathrm{eval}\,(\mathrm{map}_{\Sigma^+}g\,(b\,x)) && \{\text{by definition of }g\} \\
&= u_{g,x} && \{\text{by equation (4) with g for }h\}
\end{aligned}
$$

Blueprint Synthesis The blueprint synthesis proceeds by a straightforward syntactic analysis, similar to the one used for primitive corecursion [16]. We illustrate it with an example. Consider the definition of $\oplus$ from Sect. 3.4. Ignoring currying, the function has type $(\mathtt{nat\ stream})\,K_\oplus \to \mathtt{nat\ stream}$, with $\alpha\,K_\oplus = \alpha \times \alpha$. The term b is synthesized by processing the right-hand side of the corecursive equation for $\oplus$. After removing the syntactic sugar, we obtain the following term, highlighting the corecursive call:

$$\lambda p.\,(\mathrm{shd}\,(\mathrm{fst}\,p) + \mathrm{shd}\,(\mathrm{snd}\,p)) \lhd (\mathrm{stl}\,(\mathrm{fst}\,p) \oplus \mathrm{stl}\,(\mathrm{snd}\,p))$$

The blueprint is derived from this term by replacing the constructor guard $\lhd = \mathrm{ctor}_{\mathsf{stream}}$ and the friends with their syntactic counterparts and the corecursive call with a variable leaf:

$$b = \lambda p.\,(\mathrm{shd}\,(\mathrm{fst}\,p) + \mathrm{shd}\,(\mathrm{snd}\,p))\,\boxed{\lhd}\ \mathsf{Var}\,(\mathrm{stl}\,(\mathrm{fst}\,p), \mathrm{stl}\,(\mathrm{snd}\,p))$$

Synthesis will fail if after the indicated replacements the result does not have the desired type (here, $\mathtt{nat} \to \mathtt{nat}\,\Sigma^+$). If we omit '$(\mathrm{shd}\,(\mathrm{fst}\,p) + \mathrm{shd}\,(\mathrm{snd}\,p)) \lhd$' in the definition, the type of b becomes $\mathtt{nat} \to \mathtt{nat}\,\Sigma^*$, reflecting the lack of a guard. Another cause of failure is the presence of unfriendly operators in the call context. Once b has been produced, $\mathtt{corec}$ proves that $\oplus$ satisfies the user equation we started with.

Mixed Recursion–Corecursion If a self-call is not guarded, `corec` still gives it a chance, since it could be a terminating *recursive* call. As an example, the following definition computes all the odd numbers greater than 1 arising in the Collatz sequence:

> `corec collatz : nat → nat llist where`
> $\quad$ collatz $n = $ if $n \leq 1$ then $[]$ else if even n then $\boxed{\text{collatz } \frac{n}{2}}$ else $n \lhd$ collatz $(3 \cdot n + 1)$

The highlighted call is not guarded. Yet, it will eventually lead to a guarded call, since repeatedly halving a positive even number must at some point yield an odd number. The unguarded call yields a recursive specification of the blueprint b, which is resolved automatically by the termination prover.

By writing `corecursive` instead of `corec`, the user takes responsibility for proving termination. A manual proof was necessary for lfilter in Sect. 2.7, whose blueprint satisfies the recursion

$$b\ (P, xs) = \text{if } \forall x \in \text{set } xs.\ \neg\, P\ x \text{ then } \boxed{[]}$$
$$\text{else if } P\ (\text{lhd } xs) \text{ then lhd } xs \boxed{\lhd} \text{Var } (P, \text{ltl } xs) \text{ else } \boxed{b\ (P, \text{ltl } xs)}$$

Termination is shown by providing a suitable well-founded relation, which exists because ltl xs is closer than xs to the next element that satisfies the predicate P.

Like the corecursive calls, the recursive calls may be surrounded only by friendly operations (or by parametric operators such as 'case', 'if', and 'let'). Thus, the following specification is rejected—and rightly so, since the unfriendly stl cancels the corecursive guard that is reached when recursion terminates.

> `corec collapz : nat → nat llist where collapz` $n =$
> $\quad$ if $n = 0$ then $[]$ else if even n then $\boxed{\text{stl (collapz } \frac{n}{2})}$ else $n \lhd$ collapz $(3 \cdot n + 1)$

4.2 Registering New Friendly Operations

The command

> `corec (friend) g : J K → J where g` $x = u_{\text{g},x}$

defines g and registers it as a friend. The domain is viewed abstractly as a type constructor K applied to the codatatype J.

The command first synthesizes the blueprint $b : \text{J K} \to \text{J} \Sigma^+$, similarly to the case of plain corecursive definitions. However, this time the type Σ is not $\Sigma_{\mathcal{F}}$, but $\Sigma_{\mathcal{F}} + \text{K}$. Thus, Σ^+ mixes freely the type K with the components K_{f} of $\Sigma_{\mathcal{F}}$, which caters for *self-friendship* (as in the bfs$_2$ example from Sect. 2.4): g can be defined making use of itself as a friend (in addition to the already registered friends).

The next step is to synthesize a surface s from the blueprint b. Recall from Sect. 3.4 that a corecursively defined operator is friendly if its blueprint b can be decomposed as $s \circ \text{map}_{\text{K}} \langle \text{id}, \text{dtor} \rangle$, where $s : (\alpha \times \alpha\ \text{K}_{\text{ctor}})\ \text{K} \to \alpha\ \Sigma^+$ is parametric in α.

Once the surface s has been synthesized, proved parametric, and proved to be in the desired relationship with b, the tool invokes the STEP operation (Sect. 3.4), enriching the corecursion state with the function defined by b as a new friend, called g.

Alternatively, users can register arbitrary functions as friends:

```
friend_of_corec g : J K → J where g x = u_g,x
```

The user must then prove the equation $g\ x = u_{g,x}$. The command extracts a blueprint from it and proceeds with the surface synthesis in the same way as `corec (friend)`.

Surface Synthesis Algorithm The synthesis of the surface from the blueprint proceeds by the context-dependent replacement of some constants with terms. AmiCo performs the replacements in a logical-relation fashion, guided by type inference.

We start with $b : \mathsf{J}\ \mathsf{K} \to \mathsf{J}\ \Sigma^+$ and need to synthesize $s : (\alpha \times \alpha\ \mathsf{K_{ctor}})\ \mathsf{K} \to \alpha\ \Sigma^+$ such that s is parametric in α and $b = s \circ \mathsf{map_K}\ \langle\mathsf{id}, \mathsf{dtor}\rangle$. We traverse b recursively and collect context information about the appropriate replacements. The technical report describes the algorithm in detail. Here, we illustrate it on an example.

Consider the definition of a function that interleaves a nonempty list of streams:

```
corec (friend) inter : (nat stream) nelist → nat stream where
    inter xss = shd (hd xss) ◁ inter (tl xss ▷ stl (hd xss))
```

Here, β nelist is the type of nonempty lists with head and tail selectors $\mathsf{hd} : \beta\ \mathsf{nelist} \to \beta$ and $\mathsf{tl} : \beta\ \mathsf{nelist} \to \beta\ \mathsf{list}$ and $\triangleright : \beta\ \mathsf{list} \to \beta \to \beta\ \mathsf{nelist}$ is defined such that $xs \triangleright y$ appends y to xs. We have $\mathsf{J} = \mathsf{nat\ stream}$ and $\mathsf{K} = \mathsf{nelist}$. The blueprint is

$$b = \lambda xss.\ \boxed{\mathsf{shd}}\ (\mathsf{hd}\ xss)\ \boxed{\triangleleft}\ \mathsf{Var}\ (\mathsf{tl}\ xss \triangleright \boxed{\mathsf{stl}}\ (\mathsf{hd}\ xss))$$

From this, the tool synthesizes the surface

$$s = \lambda xss'.\ \boxed{(\mathsf{fst} \circ \mathsf{snd})}\ (\mathsf{hd}\ xss')\boxed{\triangleleft}\ \mathsf{Var}\ ((\boxed{\mathsf{map_{list}}\ \mathsf{fst} \circ \mathsf{tl}})\ xss' \triangleright \boxed{(\mathsf{snd} \circ \mathsf{snd})}\ (\mathsf{hd}\ xss'))$$

When transforming the blueprint $b : (\mathsf{nat\ stream})\ \mathsf{nelist} \to (\mathsf{nat\ stream})\ \Sigma^+$ into the surface $s : (\alpha \times (\mathsf{nat} \times \alpha))\ \mathsf{nelist} \to \alpha\ \Sigma^+$, the selectors shd and stl are replaced by suitable compositions. One of the other constants, tl, is composed with a mapping of fst. The treatment of constants is determined by their position relative to the input variables (here, xss) and by whether the input is eventually consumed by a destructor-like operator on J (here, shd and stl). Bindings can also carry consumption information—from the outer context to within their scope—as in the following variant of inter:

```
corec (friend) inter' : (nat stream) nelist → nat stream where
    inter' xss = case hd xss of x ◁ xs ⇒ x ◁ inter' (tl xss ▷ xs)
```

The case expression is syntactic sugar for a $\mathsf{case_{stream}}$ combinator. The desugared blueprint and surface constants are

$$b = \lambda xss.\ \boxed{\mathsf{case_{stream}}}\ (\mathsf{hd}\ xss)\ (\lambda x\ xs.\ x \boxed{\triangleleft}\ \mathsf{Var}\ (\mathsf{tl}\ xss \triangleright xs))$$
$$s = \lambda xss'.\ \boxed{(\mathsf{case_{prod}} \circ \mathsf{snd})}\ (\mathsf{hd}\ xss')\ (\lambda x'xs'.\ x' \boxed{\triangleleft}\ \mathsf{Var}\ ((\boxed{\mathsf{map_{list}}\ \mathsf{fst} \circ \mathsf{tl}})\ xss' \triangleright xs'))$$

The case operator for streams is processed specially, because just like shd and stl it consumes the input. The expression in the scope of the inner λ of the blueprint contains two variables—xss and xs—that have nat stream in their type. Due to the outer context, they must be treated differently: xss as an unconsumed input (which tells us to process the surrounding constant tl) and xs as a consumed input (which tells us to leave the surrounding constant $\triangleright$ unchanged). The selectors and case operators for J can also be applied indirectly, via mapping (e.g., $\mathsf{map_{nelist}}\ \mathsf{stl}\ xss$).

5 Implementation in Isabelle/HOL

The implementation of AmiCo followed the same general strategy as that of most other definitional mechanisms for Isabelle:

1. We started from an abstract formalized example consisting of a manual construction of the BASE and STEP corecursors and the corresponding reasoning principles.
2. We streamlined the formal developments, eliminating about 1000 lines of Isabelle definitions and proofs—to simplify the implementation and improve performance.
3. We formalized the new MERGE operation in the same style as BASE and STEP.
4. We developed Standard ML functions to perform the corecursor state operations for arbitrary codatatypes and friendly functions.
5. We implemented, also in Standard ML, the commands that process user specifications and interact with the corecursor state.

HOL's type system cannot express quantification over arbitrary BNFs, thus the need for ML code to repeat the corecursor derivations for each new codatatype or friend. With the foundational approach, not only the corecursors and their characteristic theorems are produced but also all the intermediate objects and lemmas, to reach the highest level of trustworthiness. Assuming the proof assistant's inference kernel is correct, bugs in our tool can lead at most to run-time failures, never to logical inconsistencies.

The code for step 4 essentially constructs the low-level types, terms, and lemma statements presented in Sect. 3 and proves the lemmas using dedicated *tactics*—ML programs that generalize the proofs from the formalization. In principle, the tactics always succeed. The code for step 5 analyses the user's specification and synthesizes blueprints and surfaces, as exemplified in Sect. 4. It reuses `primcorec`'s parsing combinators [16] for recognizing map functions and other syntactic conveniences, such as the use of λs as an alternative to $\circ$ for corecursing under $\rightarrow$, as seen in Sect. 2.1.

The archive accompanying this paper [14] contains instructions that explain where to find the code and the users' manual and how to run the code.

6 Related Work and Discussion

This work combines the safety of foundational approaches to function definitions with an expressive flavor of corecursion and mixed recursion–corecursion. It continues a program of integrating category theory insight into proof assistant technology [16–18, 67]. There is a lot of related work on corecursion and productivity, both theoretical and applied to proof assistants and functional programming languages.

Theory of (Co)recursion AmiCo incorporates category theory from many sources, notably Milius et al. [52] for corecursion up-to and Rot et al. [61] for coinduction up-to. Our earlier papers [17, 67] discuss further theoretical sources. AmiCo implements the first general, provably sound, and fully automatic method for mixing recursive and corecursive calls in function definitions. The idea of mixing recursion and corecursion appears in Bertot [11] for the stream filter, and a generalization is sketched in Bertot and Komendantskaya [13] for corecursion up to constructors. Leino's Dafny tool [46] was the first to offer such a mixture for general codatatypes, which turned out to be unsound and was subsequently restricted to the sound but limited fragment of tail recursion.

Corecursion in Other Proof Assistants Coq supports productivity by a syntactic guardedness check, based on the pioneering work of Giménez [26]. MiniAgda [2] and Agda implement a more flexible approach to productivity due to Abel et al. [3, 5], based on sized types and copatterns. Coq's guardedness check allows, in our terminology, only the constructors as friends [21]. By contrast, Agda's productivity checker is more expressive than AmiCo's, because sized types can capture more precise contracts than the "consumes at most one constructor, produces at least one constructor" criterion. For example, a Fibonacci stream definition such as $\mathsf{fib} = 0 \lhd 1 \lhd (\mathsf{fib} + \mathsf{stl}\ \mathsf{fib})$ can be made to work in Agda, but is rejected by AmiCo because stl is not a friend. As mentioned in Sect. 2.4, this flexibility comes at a price: The user must encode the productivity argument in the function's type, leading to additional proof obligations.

CIRC [50] is a theorem prover designed for automating coinduction via sound circular reasoning. It bears similarity with both Coq's Paco and our AmiCo. Its freezing operators are an antidote to what we would call the absence of friendship: Equality is no longer a congruence, hence equational reasoning is frozen at unfriendly locations.

Foundational Function Definitions AmiCo's commands and proof methods fill a gap in Isabelle/HOL's coinductive offering. They complement `codatatype`, `primcorec`, and *coinduction* [16], allowing users to define nonprimitive corecursive and mixed recursive–corecursive functions. Being foundational, our work offers a strong protection against inconsistency by reducing circular fixpoint definitions issued by the user to low-level acyclic definitions in the core logic. This approach has a long tradition.

Most systems belonging to the HOL family include a counterpart to the `primrec` command of Isabelle, which synthesizes the argument to a primitive recursor. Isabelle/HOL is the only HOL system that also supports codatatypes and `primcorec` [16]. Isabelle/ZF, for Zermelo–Fraenkel set theory, provides `(co)datatype` and `primrec` [57] commands, but no high-level mechanisms for defining corecursive functions.

For nonprimitively recursive functions over datatypes, Slind's TFL package for HOL4 and Isabelle/HOL [63] and Krauss's `function` command for Isabelle/HOL [42] are the state of the art. Krauss developed the `partial_function` command for defining monadic functions [43]. Definitional mechanisms based on the Knaster–Tarski fixpoint theorems were also developed for (co)inductive predicates [31, 57]. HOLCF, a library for domain theory, offers a `fixrec` command for defining continuous functions [35].

Our handling of friends can be seen as a round trip between a shallow and a deep embedding that resembles normalization by evaluation [9] (but starting from the shallow side). Initially, the user specification contains shallow (semantic) friends. For identifying the involved corecursion as sound, the tool reifies the friends into deep (syntactic) friends, which make up the blueprint. Then the deep friends are "reflected" back into their shallow versions by the evaluation function $\mathsf{eval} : \mathsf{J}\,\Sigma^* \to \mathsf{J}$. A similar technique is used by Myreen in HOL4 for verification and synthesis of functional programs [55].

In Agda, Coq, and Matita, the definitional mechanisms for (co)recursion are built into the system. In contrast, Lean axiomatizes only the recursor [54]. The distinguishing features of AmiCo are its dynamicity and high level of automation. The derived corecursors and coinduction principles are updated with new ones each time a friend is registered. This permits reuse both internally (resulting in lighter constructions) and at the user level (resulting in fewer proof obligations).

Code Extraction Isabelle's code generator [29] extracts Haskell code from an executable fragment of HOL, mapping HOL (co)datatypes to lazy Haskell datatypes and HOL functions to Haskell functions. Seven out of our eight case studies fall into this fragment; the extracted code is part of the archive [14]. Only the filter function on lazy lists is clearly not computable (Sect. 2.7). In particular, extraction works for Lochbihler's probabilistic calculus (Sect. 2.8) which involves the type spmf of discrete subprobability distributions. Verified data refinement in the code generator makes it possible to implement such BNFs in terms of datatypes, e.g., spmf as associative lists similar to Erwig's and Kollmansberger's PFP library [24]. Thus, we can extract code for GPVs and their operations like inlining. Lochbihler and Züst [49] used an earlier version of the calculus to implement a core of the Transport Layer Security (TLS) protocol in HOL.

Certified Lazy Programming Our tool and the examples are a first step towards a framework for friendship-based certified programming: Programs are written in the executable fragment, verified in Isabelle, and extracted to Haskell. AmiCo ensures that corecursive definitions are productive and facilitates coinductive proofs by providing strong coinduction rules. Productivity and termination of the extracted code are guaranteed if the whole program is specified in HOL exclusively with datatypes, codatatypes, recursive functions with the `function` command, and corecursive functions with `corec`, and no custom congruence rules for higher-order operators have been used. The technical report [15, Sect. 6] explains why these restrictions are necessary. If the restrictions are met, the program clearly lies within the executable fragment and the code extracted from the definitions yields the higher-order rewrite system which the termination prover and AmiCo have checked. In particular, these restrictions exclude the noncomputable filter function on lazy lists (Sect. 2.7), with the test $\forall n \in \text{set } xs. \neg P \, n$.

A challenge will be to extend these guarantees to Isabelle's modular architecture. Having been designed with only partial correctness in mind, the code extractor can be customized to execute arbitrary (proved) equations—which can easily break productivity and termination. A similar issue occurs with `friend_of_corec`, which cares only about semantic properties of the friend to be. For example, we can specify the identity function id on streams by id $(x \lhd y \lhd xs) = x \lhd y \lhd xs$ and register it as a friend with the derived equation id $x = \text{shd } x \lhd \text{stl } x$. Consequently, AmiCo accepts the definition natsFrom $n = n \lhd \text{id } (\text{natsFrom } (n+1))$, but the extracted Haskell code diverges. To avoid these problems, we would have to (re)check productivity and termination on the equations used for extraction. In this scenario, AmiCo can be used to distinguish recursive from corecursive calls in a set of (co)recursive equations, and synthesize sufficient conditions for the function being productive and the recursion terminating, and automatically prove them (using Isabelle's parametricity [36] and termination provers [20]).

AmiCo beyond Higher-Order Logic The techniques implemented in our tool are applicable beyond Isabelle/HOL. In principle, nothing stands in the way of AgdamiCo, AmiCoq, or MatitamiCo. Danielsson [22] and Thibodeau et al. [65] showed that similar approaches work in type theory; what is missing is a tool design and implementation. AmiCo relies on parametricity, which is now understood for dependent types [10].

In Agda, parametricity could be encoded with sized types, and AgdamiCo could be a foundational tool that automatically adds suitable sized types for justifying the def-

inition and erases them from the end product. Coq includes a parametricity-tracking tool [40] that could form the basis of AmiCoq. The Paco library by Hur et al. [37] facilitates coinductive proofs based on parameterized coinduction [53, 70]. Recent work by Pous [59] includes a framework to combine proofs by induction and coinduction. An AmiCoq would catch up on the corecursion definition front, going beyond what is possible with the *cofix* tactic [21]. On the proof front, AmiCoq would provide a substantial entry into Paco's knowledge base: For any codatatype J with destructor $dtor : J \to J\,K$, all registered friends are, in Paco's terminology, respectful up-to functions for the monotonic operator $\lambda r\,x\,y.\ rel_K\,r\,(dtor\,x)\,(dtor\,y)$, whose greatest fixpoint is the equality on J.

A more lightweight application of our methodology would be an AmiCo for Haskell or for more specialized languages such as CoCaml [38]. In these languages, parametricity is ensured by the computational model. An automatic tool that embodies AmiCo's principles could analyze a Haskell program and prove it total. For CoCaml, which is total, a tool could offer more flexibility when writing corecursive programs.

Surface Synthesis beyond Corecursion The notion of extracting a parametric component with suitable properties can be useful in other contexts than corecursion. In the programming-by-examples paradigm [28], one needs to choose between several synthesized programs whose behavior matches a set of input–output instances. These criteria tend to prefer programs that are highly parametric. A notion of degree of parametricity does not exist in the literature but could be expressed as the size of a parametric surface, for a suitable notion of surface, where $\langle id, dtor \rangle$ is replaced by domain specific functions and fst by their left inverses.

References

1. Abbott, M., Altenkirch, T., Ghani, N.: Containers: Constructing strictly positive types. Theor. Comput. Sci. 342(1), 3–27 (2005)
2. Abel, A.: MiniAgda: Integrating sized and dependent types. In: Bove, A., Komendantskaya, E., Niqui, M. (eds.) PAR 2010. EPTCS, vol. 43, pp. 14–28 (2010)
3. Abel, A.: Compositional coinduction with sized types. In: Hasuo, I. (ed.) CMCS 2016. LNCS, vol. 9608, pp. 5–10. Springer (2016)
4. Abel, A., Pientka, B.: Well-founded recursion with copatterns and sized types. J. Funct. Program. 26, e2 (2016)

5. Abel, A., Pientka, B., Thibodeau, D., Setzer, A.: Copatterns: Programming infinite structures by observations. In: Giacobazzi, R., Cousot, R. (eds.) POPL 2013. pp. 27–38. ACM (2013)

6. Adams, M.: Introducing HOL Zero (extended abstract). In: Fukuda, K., van der Hoeven, J., Joswig, M., Takayama, N. (eds.) ICMS 2010. LNCS, vol. 6327, pp. 142–143. Springer (2010)

7. Asperti, A., Ricciotti, W., Coen, C.S., Tassi, E.: The Matita interactive theorem prover. In: Bjørner, N., Sofronie-Stokkermans, V. (eds.) CADE-23. LNCS, vol. 6803, pp. 64–69. Springer (2011)

8. Atkey, R., McBride, C.: Productive coprogramming with guarded recursion. In: Morrisett, G., Uustalu, T. (eds.) ICFP 2013. pp. 197–208. ACM (2013)

9. Berger, U., Schwichtenberg, H.: An inverse of the evaluation functional for typed lambda-calculus. In: LICS '91. pp. 203–211. IEEE Computer Society (1991)

10. Bernardy, J.P., Jansson, P., Paterson, R.: Proofs for free: Parametricity for dependent types. J. Funct. Program. 22(2), 107–152 (2012)

11. Bertot, Y.: Filters on coinductive streams, an application to Eratosthenes' sieve. In: Urzyczyn, P. (ed.) TLCA 2005. LNCS, vol. 3461, pp. 102–115. Springer (2005)

12. Bertot, Y., Casteran, P.: Interactive Theorem Proving and Program Development—Coq'Art: The Calculus of Inductive Constructions. Texts in Theoretical Computer Science, Springer (2004)

13. Bertot, Y., Komendantskaya, E.: Inductive and coinductive components of corecursive functions in Coq. Electr. Notes Theor. Comput. Sci. 203(5), 25–47 (2008)

14. Blanchette, J.C., Bouzy, A., Lochbihler, A., Popescu, A., Traytel, D.: Archive associated with this paper. `http://matryoshka.gforge.inria.fr/pubs/amico_material.tar.gz`

15. Blanchette, J.C., Bouzy, A., Lochbihler, A., Popescu, A., Traytel, D.: Friends with benefits: Implementing corecursion in foundational proof assistants. Tech. rep. (2017), `http://matryoshka.gforge.inria.fr/pubs/amico_report.pdf`

16. Blanchette, J.C., Hölzl, J., Lochbihler, A., Panny, L., Popescu, A., Traytel, D.: Truly modular (co)datatypes for Isabelle/HOL. In: Klein, G., Gamboa, R. (eds.) ITP 2014. LNCS, vol. 8558, pp. 93–110. Springer (2014)

17. Blanchette, J.C., Popescu, A., Traytel, D.: Foundational extensible corecursion: A proof assistant perspective. In: Fisher, K., Reppy, J.H. (eds.) ICFP 2015. pp. 192–204. ACM (2015)

18. Blanchette, J.C., Popescu, A., Traytel, D.: Witnessing (co)datatypes. In: Vitek, J. (ed.) ESOP 2015. LNCS, vol. 9032, pp. 359–382. Springer (2015)

19. Bove, A., Dybjer, P., Norell, U.: A brief overview of Agda—A functional language with dependent types. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLs 2009. LNCS, vol. 5674, pp. 73–78. Springer (2009)

20. Bulwahn, L., Krauss, A., Nipkow, T.: Finding lexicographic orders for termination proofs in Isabelle/HOL. In: Schneider, K., Brandt, J. (eds.) TPHOLs 2007. LNCS, vol. 4732, pp. 38–53. Springer (2007)

21. Chlipala, A.: Certified Programming with Dependent Types—A Pragmatic Introduction to the Coq Proof Assistant. MIT Press (2013)

22. Danielsson, N.A.: Beating the productivity checker using embedded languages. In: Bove, A., Komendantskaya, E., Niqui, M. (eds.) PAR 2010. EPTCS, vol. 43, pp. 29–48 (2010)

23. Dijkstra, E.W.: An exercise for Dr. R. M. Burstall. In: Selected Writings on Computing: A Personal Perspective, pp. 215–216. Texts and Monographs in Computer Science, Springer (1982)

24. Erwig, M., Kollmansberger, S.: Probabilistic functional programming in Haskell. J. Funct. Programm. 16(1), 21–34 (2006)

25. Gammie, P., Lochbihler, A.: The Stern–Brocot tree. Archive of Formal Proofs (2015), `https://www.isa-afp.org/entries/Stern_Brocot.shtml`

26. Giménez, E.: Codifying guarded definitions with recursive schemes. In: Dybjer, P., Nordström, B., Smith, J.M. (eds.) TYPES '94. LNCS, vol. 996, pp. 39–59. Springer (1995)

27. Gordon, M.J.C., Melham, T.F. (eds.): Introduction to HOL: A Theorem Proving Environment for Higher Order Logic. Cambridge University Press (1993)

28. Gulwani, S.: Programming by examples—and its applications in data wrangling. In: Dependable Software Systems Engineering, NATO Science for Peace and Security Series—D: Information and Communication Security, vol. 45, pp. 137–158. IOS Press (2016)

29. Haftmann, F., Nipkow, T.: Code generation via higher-order rewrite systems. In: Blume, M., Kobayashi, N., Vidal, G. (eds.) FLOPS 2010. LNCS, vol. 6009, pp. 103–117. Springer (2010)

30. Hagino, T.: A Categorical Programming Language. Ph.D. thesis, University of Edinburgh (1987)

31. Harrison, J.: Inductive definitions: Automation and application. In: Schubert, E.T., Windley, P.J., Alves-Foss, J. (eds.) TPHOLs '95. LNCS, vol. 971, pp. 200–213. Springer (1995)

32. Harrison, J.: HOL Light: An overview. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLs 2009. LNCS, vol. 5674, pp. 60–66. Springer (2009)

33. Hinze, R.: The Bird tree. J. Func. Programm. 19(5), 491–508 (2009)

34. Homeier, P.V.: The HOL-Omega logic. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLs 2009. LNCS, vol. 5674, pp. 244–259. Springer (2009)

35. Huffman, B.: HOLCF '11: A Definitional Domain Theory for Verifying Functional Programs. Ph.D. thesis, Portland State University (2012)

36. Huffman, B., Kunčar, O.: Lifting and Transfer: A modular design for quotients in Isabelle/HOL. In: Gonthier, G., Norrish, M. (eds.) CPP 2013. LNCS, vol. 8307, pp. 131–146. Springer (2013)

37. Hur, C.K., Neis, G., Dreyer, D., Vafeiadis, V.: The power of parameterization in coinductive proof. In: Giacobazzi, R., Cousot, R. (eds.) POPL 2013. pp. 193–206. ACM (2013)

38. Jeannin, J., Kozen, D., Silva, A.: Language constructs for non-well-founded computation. In: Felleisen, M., Gardner, P. (eds.) ESOP 2013. LNCS, vol. 7792, pp. 61–80. Springer (2013)

39. Jones, G., Gibbons, J.: Linear-time breadth-first tree algorithms: An exercise in the arithmetic of folds and zips. Tech. Rep. 71, C.S. Dept., University of Auckland (1993)

40. Keller, C., Lasson, M.: Parametricity in an impredicative sort. In: Cégielski, P., Durand, A. (eds.) CSL 2012. LIPIcs, vol. 16, pp. 381–395. Schloss Dagstuhl—Leibniz-Zentrum für Informatik (2012)

41. Knuth, D.E., Morris, J.H., Pratt, V.R.: Fast pattern matching in strings. SIAM J. Comput. 6(2), 323–350 (1977)

42. Krauss, A.: Partial recursive functions in higher-order logic. In: Furbach, U., Shankar, N. (eds.) IJCAR 2006. LNCS, vol. 4130, pp. 589–603. Springer (2006)

43. Krauss, A.: Recursive definitions of monadic functions. In: Bove, A., Komendantskaya, E., Niqui, M. (eds.) PAR 2010. EPTCS, vol. 43, pp. 1–13 (2010)

44. van Laarhoven, T.: Knuth–Morris–Pratt in Haskell. http://www.twanvl.nl/blog/haskell/Knuth-Morris-Pratt-in-Haskell (2007)

45. Leino, K.R.M.: Automating theorem proving with SMT. In: Blazy, S., Paulin-Mohring, C., Pichardie, D. (eds.) ITP 2013. LNCS, vol. 7998, pp. 2–16. Springer (2013)

46. Leino, K.R.M., Moskal, M.: Co-induction simply: Automatic co-inductive proofs in a program verifier. In: Jones, C.B., Pihlajasaari, P., Sun, J. (eds.) FM 2014. LNCS, vol. 8442, pp. 382–398. Springer (2014)

47. Lochbihler, A.: Probabilistic functions and cryptographic oracles in higher-order logic. In: Thiemann, P. (ed.) ESOP 2016. LNCS, vol. 9632, pp. 503–531. Springer (2016)

48. Lochbihler, A., Hölzl, J.: Recursive functions on lazy lists via domains and topologies. In: Klein, G., Gamboa, R. (eds.) ITP 2014. LNCS, vol. 8558, pp. 341–357. Springer (2014)

49. Lochbihler, A., Züst, M.: Programming TLS in Isabelle/HOL. Isabelle Workshop 2014 (2014), `https://www.ethz.ch/content/dam/ethz/special-interest/infk/inst-infsec/information-security-group-dam/research/publications/pub2014/lochbihler14iw.pdf`

50. Lucanu, D., Goriac, E.I., Caltais, G., Roşu, G.: CIRC: A behavioral verification tool based on circular coinduction. In: Kurz, A., Lenisa, M., Tarlecki, A. (eds.) CALCO 2009. LNCS, vol. 5728, pp. 433–442. Springer (2009)

51. Milius, S.: Completely iterative algebras and completely iterative monads. Inf. Comput. 196(1), 1–41 (2005)

52. Milius, S., Moss, L.S., Schwencke, D.: Abstract GSOS rules and a modular treatment of recursive definitions. Log. Meth. Comput. Sci. 9(3) (2013)

53. Moss, L.S.: Parametric corecursion. Theor. Comput. Sci. 260(1-2), 139–163 (2001)

54. de Moura, L., Kong, S., Avigad, J., van Doorn, F., von Raumer, J.: The Lean theorem prover (system description). In: Felty, A.P., Middeldorp, A. (eds.) CADE-25. LNCS, vol. 9195, pp. 378–388. Springer (2015)

55. Myreen, M.O.: Functional programs: Conversions between deep and shallow embeddings. In: Beringer, L., Felty, A. (eds.) ITP 2012. pp. 412–417. LNCS, Springer (2012)

56. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL: A Proof Assistant for Higher-Order Logic. Springer (2002)

57. Paulson, L.C.: A fixedpoint approach to implementing (co)inductive definitions. In: Bundy, A. (ed.) CADE-12. LNCS, vol. 814, pp. 148–161. Springer (1994)

58. Paulson, L.C.: Mechanizing coinduction and corecursion in higher-order logic. J. Log. Comput. 7(2), 175–204 (1997)

59. Pous, D.: Coinduction all the way up. In: Grohe, M., Koskinen, E., Shankar, N. (eds.) LICS 2016. pp. 307–316. ACM (2016)

60. Reynolds, J.C.: Types, abstraction and parametric polymorphism. In: Mason, R.E.A. (ed.) IFIP '83. pp. 513–523. North-Holland/IFIP (1983)

61. Rot, J., Bonsangue, M.M., Rutten, J.J.M.M.: Coalgebraic bisimulation-up-to. In: van Emde Boas, P., Groen, F.C.A., Italiano, G.F., Nawrocki, J.R., Sack, H. (eds.) SOFSEM 2013. LNCS, vol. 7741, pp. 369–381. Springer (2013)

62. Rutten, J.J.M.M.: Automata and coinduction (an exercise in coalgebra). In: Sangiorgi, D., de Simone, R. (eds.) CONCUR '98. LNCS, vol. 1466, pp. 194–218. Springer (1998)

63. Slind, K.: Function definition in higher-order logic. In: von Wright, J., Grundy, J., Harrison, J. (eds.) TPHOLs '96. LNCS, vol. 1125, pp. 381–397. Springer (1996)

64. Slind, K., Norrish, M.: A brief overview of HOL4. In: Mohamed, O.A., Muñoz, C., Tahar, S. (eds.) TPHOLs 2008. LNCS, vol. 5170, pp. 28–32. Springer (2008)

65. Thibodeau, D., Cave, A., Pientka, B.: Indexed codata types. In: Sumii, E. (ed.) ICFP 2016. ACM (2016)

66. Traytel, D.: Formal languages, formally and coinductively. In: Kesner, D., Pientka, B. (eds.) FSCD 2016. LIPIcs, vol. 52, pp. 31:1–17. Schloss Dagstuhl—Leibniz-Zentrum für Informatik (2016)

67. Traytel, D., Popescu, A., Blanchette, J.C.: Foundational, compositional (co)datatypes for higher-order logic: Category theory applied to theorem proving. In: LICS 2012, pp. 596–605. IEEE Computer Society (2012)

68. Turner, D.A.: Elementary strong functional programming. In: Hartel, P.H., Plasmeijer, M.J. (eds.) FPLE '95. LNCS, vol. 1022, pp. 1–13. Springer (1995)

69. Wadler, P.: Theorems for free! In: Stoy, J.E. (ed.) FPCA 1989. pp. 347–359. ACM (1989)

70. Winskel, G.: A note on model checking the modal v-calculus. Theor. Comput. Sci. 83(1), 157–167 (1991)