A Fixedpoint Approach to (Co)Inductive and (Co)Datatype Definitions*

Lawrence C. Paulson

`lcp@cl.cam.ac.uk`

Computer Laboratory, University of Cambridge, England

21 June 2010

Abstract

This paper presents a fixedpoint approach to inductive definitions. Instead of using a syntactic test such as "strictly positive," the approach lets definitions involve any operators that have been proved monotone. It is conceptually simple, which has allowed the easy implementation of mutual recursion and iterated definitions. It also handles coinductive definitions: simply replace the least fixedpoint by a greatest fixedpoint.

The method has been implemented in two of Isabelle's logics, ZF set theory and higher-order logic. It should be applicable to any logic in which the Knaster-Tarski theorem can be proved. Examples include lists of n elements, the accessible part of a relation and the set of primitive recursive functions. One example of a coinductive definition is bisimulations for lazy lists. Recursive datatypes are examined in detail, as well as one example of a **codatatype**: lazy lists.

The Isabelle package has been applied in several large case studies, including two proofs of the Church-Rosser theorem and a coinductive proof of semantic consistency. The package can be trusted because it proves theorems from definitions, instead of asserting desired properties as axioms.

Contents

1 Introduction

Several theorem provers provide commands for formalizing recursive data structures, like lists and trees. Robin Milner implemented one of the first of these, for Edinburgh LCF [15]. Given a description of the desired data structure, Milner's package formulated appropriate definitions and proved the characteristic theorems. Similar is Melham's recursive type package for the Cambridge HOL system [14]. Such data structures are called **datatypes** below, by analogy with datatype declarations in Standard ML. Some logics take datatypes as primitive; consider Boyer and Moore's shell principle [4] and the Coq type theory [21].

A datatype is but one example of an **inductive definition**. Such a definition [2] specifies the least set R **closed under** given rules: applying a rule to elements of R yields a result within R. Inductive definitions have many applications. The collection of theorems in a logic is inductively defined. A structural operational semantics [12] is an inductive definition of a reduction or evaluation relation on programs. A few theorem provers provide commands for formalizing inductive definitions; these include Coq [21] and again the HOL system [5].

The dual notion is that of a **coinductive definition**. Such a definition specifies the greatest set R **consistent with** given rules: every element of R can be seen as arising by applying a rule to elements of R. Important examples include using bisimulation relations to formalize equivalence of processes [16] or lazy functional programs [1]. Other examples include lazy lists and other infinite data structures; these are called **codatatypes** below.

Not all inductive definitions are meaningful. **Monotone** inductive definitions are a large, well-behaved class. Monotonicity can be enforced by syntactic conditions such as "strictly positive," but this could lead to monotone definitions being rejected on the grounds of their syntactic form. More flexible is to formalize monotonicity within the logic and allow users to prove it.

This paper describes a package based on a fixedpoint approach. Least fixedpoints yield inductive definitions; greatest fixedpoints yield coinductive definitions. Most of the discussion below applies equally to inductive and coinductive definitions, and most of the code is shared.

The package supports mutual recursion and infinitely-branching datatypes and codatatypes. It allows use of any operators that have been proved monotone, thus accepting all provably monotone inductive definitions, including iterated definitions.

The package has been implemented in Isabelle [28, 24] using ZF set theory [23, 25]; part of it has since been ported to Isabelle/HOL (higher-order

logic). The recursion equations are specified as introduction rules for the mutually recursive sets. The package transforms these rules into a mapping over sets, and attempts to prove that the mapping is monotonic and well-typed. If successful, the package makes fixedpoint definitions and proves the introduction, elimination and (co)induction rules. Users invoke the package by making simple declarations in Isabelle theory files.

Most datatype packages equip the new datatype with some means of expressing recursive functions. This is the main omission from my package. Its fixedpoint operators define only recursive sets. The Isabelle/ZF theory provides well-founded recursion [25], which is harder to use than structural recursion but considerably more general. Slind [33] has written a package to automate the definition of well-founded recursive functions in Isabelle/HOL.

Outline. Section 2 introduces the least and greatest fixedpoint operators. Section 3 discusses the form of introduction rules, mutual recursion and other points common to inductive and coinductive definitions. Section 4 discusses induction and coinduction rules separately. Section 5 presents several examples, including a coinductive definition. Section 6 describes datatype definitions. Section 7 presents related work. Section 8 draws brief conclusions. The appendices are simple user's manuals for this Isabelle package.

Most of the definitions and theorems shown below have been generated by the package. I have renamed some variables to improve readability.

2   Fixedpoint operators

In set theory, the least and greatest fixedpoint operators are defined as follows:

$$\mathtt{lfp}(D, h) \;\equiv\; \bigcap \{X \subseteq D \,.\, h(X) \subseteq X\}$$
$$\mathtt{gfp}(D, h) \;\equiv\; \bigcup \{X \subseteq D \,.\, X \subseteq h(X)\}$$

Let D be a set. Say that h is **bounded by** D if $h(D) \subseteq D$, and **monotone below** D if $h(A) \subseteq h(B)$ for all A and B such that $A \subseteq B \subseteq D$. If h is bounded by D and monotone then both operators yield fixedpoints:

$$\mathtt{lfp}(D, h) \;=\; h(\mathtt{lfp}(D, h))$$
$$\mathtt{gfp}(D, h) \;=\; h(\mathtt{gfp}(D, h))$$

These equations are instances of the Knaster-Tarski theorem, which states that every monotonic function over a complete lattice has a fixedpoint [6]. It

is obvious from their definitions that `lfp` must be the least fixedpoint, and `gfp` the greatest.

This fixedpoint theory is simple. The Knaster-Tarski theorem is easy to prove. Showing monotonicity of h is trivial, in typical cases. We must also exhibit a bounding set D for h. Frequently this is trivial, as when a set of theorems is (co)inductively defined over some previously existing set of formulæ. Isabelle/ZF provides suitable bounding sets for infinitely-branching (co)datatype definitions; see §6.1. Bounding sets are also called **domains**.

The powerset operator is monotone, but by Cantor's theorem there is no set A such that $A = \mathcal{P}(A)$. We cannot put $A = \texttt{lfp}(D, \mathcal{P})$ because there is no suitable domain D. But §5.5 demonstrates that $\mathcal{P}$ is still useful in inductive definitions.

3   Elements of an inductive or coinductive definition

Consider a (co)inductive definition of the sets R_1, ..., R_n, in mutual recursion. They will be constructed from domains D_1, ..., D_n, respectively. The construction yields not $R_i \subseteq D_i$ but $R_i \subseteq D_1 + \cdots + D_n$, where R_i is contained in the image of D_i under an injection. Reasons for this are discussed elsewhere [25, §4.5].

The definition may involve arbitrary parameters $\vec{p} = p_1$, ..., p_k. Each recursive set then has the form $R_i(\vec{p})$. The parameters must be identical every time they occur within a definition. This would appear to be a serious restriction compared with other systems such as Coq [21]. For instance, we cannot define the lists of n elements as the set $\texttt{listn}(A, n)$ using rules where the parameter n varies. Section 5.2 describes how to express this set using the inductive definition package.

To avoid clutter below, the recursive sets are shown as simply R_i instead of $R_i(\vec{p})$.

3.1   The form of the introduction rules

The body of the definition consists of the desired introduction rules. The conclusion of each rule must have the form $t \in R_i$, where t is any term. Premises typically have the same form, but they can have the more general form $t \in M(R_i)$ or express arbitrary side-conditions.

The premise $t \in M(R_i)$ is permitted if M is a monotonic operator on

sets, satisfying the rule

$$\frac{A \subseteq B}{M(A) \subseteq M(B)}$$

The user must supply the package with monotonicity rules for all such premises.

The ability to introduce new monotone operators makes the approach flexible. A suitable choice of M and t can express a lot. The powerset operator $\mathcal{P}$ is monotone, and the premise $t \in \mathcal{P}(R)$ expresses $t \subseteq R$; see §5.5 for an example. The *list of* operator is monotone, as is easily proved by induction. The premise $t \in \texttt{list}(R)$ avoids having to encode the effect of $\texttt{list}(R)$ using mutual recursion; see §5.6 and also my earlier paper [25, §4.4].

Introduction rules may also contain **side-conditions**. These are premises consisting of arbitrary formulæ not mentioning the recursive sets. Side-conditions typically involve type-checking. One example is the premise $a \in A$ in the following rule from the definition of lists:

$$\frac{a \in A \quad l \in \texttt{list}(A)}{\texttt{Cons}(a, l) \in \texttt{list}(A)}$$

3.2   The fixedpoint definitions

The package translates the list of desired introduction rules into a fixedpoint definition. Consider, as a running example, the finite powerset operator $\texttt{Fin}(A)$: the set of all finite subsets of A. It can be defined as the least set closed under the rules

$$\emptyset \in \texttt{Fin}(A) \qquad \frac{a \in A \quad b \in \texttt{Fin}(A)}{\{a\} \cup b \in \texttt{Fin}(A)}$$

The domain in a (co)inductive definition must be some existing set closed under the rules. A suitable domain for $\texttt{Fin}(A)$ is $\mathcal{P}(A)$, the set of all subsets of A. The package generates the definition

$$\texttt{Fin}(A) \equiv \texttt{lfp}(\mathcal{P}(A), \ \lambda X \,.\, \{z \in \mathcal{P}(A).\, z = \emptyset \,\vee \\ (\exists a\, b \,.\, z = \{a\} \cup b \wedge a \in A \wedge b \in X)\})$$

The contribution of each rule to the definition of $\texttt{Fin}(A)$ should be obvious. A coinductive definition is similar but uses $\texttt{gfp}$ instead of $\texttt{lfp}$.

The package must prove that the fixedpoint operator is applied to a monotonic function. If the introduction rules have the form described above, and if

the package is supplied a monotonicity theorem for every $t \in M(R_i)$ premise, then this proof is trivial.[1]

The package returns its result as an ML structure, which consists of named components; we may regard it as a record. The result structure contains the definitions of the recursive sets as a theorem list called `defs`. It also contains some theorems; `dom_subset` is an inclusion such as $\texttt{Fin}(A) \subseteq \mathcal{P}(A)$, while `bnd_mono` asserts that the fixedpoint definition is monotonic.

Internally the package uses the theorem `unfold`, a fixedpoint equation such as

$$\texttt{Fin}(A) = \{z \in \mathcal{P}(A).\, z = \emptyset \vee \\ (\exists a\, b\,.\, z = \{a\} \cup b \wedge a \in A \wedge b \in \texttt{Fin}(A))\}$$

In order to save space, this theorem is not exported.

3.3 Mutual recursion

In a mutually recursive definition, the domain of the fixedpoint construction is the disjoint sum of the domain D_i of each R_i, for $i = 1, \ldots, n$. The package uses the injections of the binary disjoint sum, typically `Inl` and `Inr`, to express injections $h_{1n}, \ldots, h_{nn}$ for the n-ary disjoint sum $D_1 + \cdots + D_n$.

As discussed elsewhere [25, §4.5], Isabelle/ZF defines the operator `Part` to support mutual recursion. The set $\texttt{Part}(A, h)$ contains those elements of A having the form $h(z)$:

$$\texttt{Part}(A, h) \equiv \{x \in A\,.\,\exists z\,.\, x = h(z)\}.$$

For mutually recursive sets $R_1, \ldots, R_n$ with $n > 1$, the package makes $n + 1$ definitions. The first defines a set R using a fixedpoint operator. The remaining n definitions have the form

$$R_i \equiv \texttt{Part}(R, h_{in}), \qquad i = 1, \ldots, n.$$

It follows that $R = R_1 \cup \cdots \cup R_n$, where the R_i are pairwise disjoint.

3.4 Proving the introduction rules

The user supplies the package with the desired form of the introduction rules. Once it has derived the theorem `unfold`, it attempts to prove those rules.

---

[1]Due to the presence of logical connectives in the fixedpoint's body, the monotonicity proof requires some unusual rules. These state that the connectives $\wedge$, $\vee$ and $\exists$ preserve monotonicity with respect to the partial ordering on unary predicates given by $P \sqsubseteq Q$ if and only if $\forall x\,.\, P(x) \rightarrow Q(x)$.

From the user's point of view, this is the trickiest stage; the proofs often fail. The task is to show that the domain $D_1 + \cdots + D_n$ of the combined set $R_1 \cup \cdots \cup R_n$ is closed under all the introduction rules. This essentially involves replacing each R_i by $D_1 + \cdots + D_n$ in each of the introduction rules and attempting to prove the result.

Consider the $\mathtt{Fin}(A)$ example. After substituting $\mathcal{P}(A)$ for $\mathtt{Fin}(A)$ in the rules, the package must prove

$$\emptyset \in \mathcal{P}(A) \qquad \frac{a \in A \quad b \in \mathcal{P}(A)}{\{a\} \cup b \in \mathcal{P}(A)}$$

Such proofs can be regarded as type-checking the definition.[2] The user supplies the package with type-checking rules to apply. Usually these are general purpose rules from the ZF theory. They could however be rules specifically proved for a particular inductive definition; sometimes this is the easiest way to get the definition through!

The result structure contains the introduction rules as the theorem list `intrs`.

3.5 The case analysis rule

The elimination rule, called `elim`, performs case analysis. It is a simple consequence of `unfold`. There is one case for each introduction rule. If $x \in \mathtt{Fin}(A)$ then either $x = \emptyset$ or else $x = \{a\} \cup b$ for some $a \in A$ and $b \in \mathtt{Fin}(A)$. Formally, the elimination rule for $\mathtt{Fin}(A)$ is written

$$\frac{x \in \mathtt{Fin}(A) \qquad \begin{array}{c} [x = \emptyset] \\ \vdots \\ Q \end{array} \qquad \begin{array}{c} [x = \{a\} \cup b \quad a \in A \quad b \in \mathtt{Fin}(A)]_{a,b} \\ \vdots \\ Q \end{array}}{Q}$$

The subscripted variables a and b above the third premise are eigenvariables, subject to the usual "not free in ..." proviso.

4 Induction and coinduction rules

Here we must consider inductive and coinductive definitions separately. For an inductive definition, the package returns an induction rule derived directly from the properties of least fixedpoints, as well as a modified rule for mutual recursion. For a coinductive definition, the package returns a basic coinduction rule.

---

[2]The Isabelle/HOL version does not require these proofs, as HOL has implicit type-checking.

4.1 The basic induction rule

The basic rule, called `induct`, is appropriate in most situations. For inductive definitions, it is strong rule induction [5]; for datatype definitions (see below), it is just structural induction.

The induction rule for an inductively defined set R has the form described below. For the time being, assume that R's domain is not a Cartesian product; inductively defined relations are treated slightly differently.

The major premise is $x \in R$. There is a minor premise for each introduction rule:

- If the introduction rule concludes $t \in R_i$, then the minor premise is $P(t)$.

- The minor premise's eigenvariables are precisely the introduction rule's free variables that are not parameters of R. For instance, the eigenvariables in the $\mathtt{Fin}(A)$ rule below are a and b, but not A.

- If the introduction rule has a premise $t \in R_i$, then the minor premise discharges the assumption $t \in R_i$ and the induction hypothesis $P(t)$. If the introduction rule has a premise $t \in M(R_i)$ then the minor premise discharges the single assumption

$$t \in M(\{z \in R_i \,.\, P(z)\}).$$

  Because M is monotonic, this assumption implies $t \in M(R_i)$. The occurrence of P gives the effect of an induction hypothesis, which may be exploited by appealing to properties of M.

The induction rule for $\mathtt{Fin}(A)$ resembles the elimination rule shown above, but includes an induction hypothesis:

$$\frac{x \in \mathtt{Fin}(A) \quad P(\emptyset) \qquad \begin{array}{c} [a \in A \quad b \in \mathtt{Fin}(A) \quad P(b)]_{a,b} \\ \vdots \\ P(\{a\} \cup b) \end{array}}{P(x)}$$

Stronger induction rules often suggest themselves. We can derive a rule for $\mathtt{Fin}(A)$ whose third premise discharges the extra assumption $a \notin b$. The package provides rules for mutual induction and inductive relations. The Isabelle/ZF theory also supports well-founded induction and recursion over datatypes, by reasoning about the **rank** of a set [25, §3.4].

4.2   Modified induction rules

If the domain of R is a Cartesian product $A_1 \times \cdots \times A_m$ (however nested), then the corresponding predicate P_i takes m arguments. The major premise becomes $\langle z_1, \ldots, z_m \rangle \in R$ instead of $x \in R$; the conclusion becomes $P(z_1, \ldots, z_m)$. This simplifies reasoning about inductively defined relations, eliminating the need to express properties of z_1, ..., z_m as properties of the tuple $\langle z_1, \ldots, z_m \rangle$. Occasionally it may require you to split up the induction variable using `SigmaE` and `dom_subset`, especially if the constant `split` appears in the rule.

The mutual induction rule is called `mutual_induct`. It differs from the basic rule in two respects:

- Instead of a single predicate P, it uses n predicates P_1, ..., P_n: one for each recursive set.

- There is no major premise such as $x \in R_i$. Instead, the conclusion refers to all the recursive sets:

$$(\forall z \,.\, z \in R_1 \to P_1(z)) \wedge \cdots \wedge (\forall z \,.\, z \in R_n \to P_n(z))$$

  Proving the premises establishes $P_i(z)$ for $z \in R_i$ and $i = 1$, ..., n.

If the domain of some R_i is a Cartesian product, then the mutual induction rule is modified accordingly. The predicates are made to take m separate arguments instead of a tuple, and the quantification in the conclusion is over the separate variables z_1, ..., z_m.

4.3   Coinduction

A coinductive definition yields a primitive coinduction rule, with no refinements such as those for the induction rules. (Experience may suggest refinements later.) Consider the codatatype of lazy lists as an example. For suitable definitions of `LNil` and `LCons`, lazy lists may be defined as the greatest set consistent with the rules

$$\text{LNil} \in \text{llist}(A) \qquad \frac{a \in A \quad l \in \text{llist}(A)}{\text{LCons}(a, l) \in \text{llist}(A)} \; (-)$$

The $(-)$ tag stresses that this is a coinductive definition. A suitable domain for `llist`(A) is `quniv`(A); this set is closed under the variant forms of sum and product that are used to represent non-well-founded data structures (see §6.1).

The package derives an `unfold` theorem similar to that for `Fin`(A). Then it proves the theorem `coinduct`, which expresses that `llist`(A) is the greatest solution to this equation contained in `quniv`(A):

$$
\cfrac{x \in X \quad X \subseteq \mathtt{quniv}(A) \quad \begin{array}{c} [z \in X]_z \\ \vdots \\ z = \mathtt{LNil} \vee (\exists a\, l.\, z = \mathtt{LCons}(a, l) \wedge a \in A\, \wedge \\ l \in X \cup \mathtt{llist}(A)) \end{array}}{x \in \mathtt{llist}(A)}
$$

This rule complements the introduction rules; it provides a means of showing $x \in \mathtt{llist}(A)$ when x is infinite. For instance, if $x = \mathtt{LCons}(0, x)$ then applying the rule with $X = \{x\}$ proves $x \in \mathtt{llist}(\mathtt{nat})$. (Here `nat` is the set of natural numbers.)

Having $X \cup \mathtt{llist}(A)$ instead of simply X in the third premise above represents a slight strengthening of the greatest fixedpoint property. I discuss several forms of coinduction rules elsewhere [26].

The clumsy form of the third premise makes the rule hard to use, especially in large definitions. Probably a constant should be declared to abbreviate the large disjunction, and rules derived to allow proving the separate disjuncts.

5 Examples of inductive and coinductive definitions

This section presents several examples from the literature: the finite powerset operator, lists of n elements, bisimulations on lazy lists, the well-founded part of a relation, and the primitive recursive functions.

5.1 The finite powerset operator

This operator has been discussed extensively above. Here is the corresponding invocation in an Isabelle theory file. Note that `cons`(a, b) abbreviates $\{a\} \cup b$ in Isabelle/ZF.

```
Finite = Arith +
consts      Fin :: i=>i
inductive
  domains   "Fin(A)" <= "Pow(A)"
  intrs
    emptyI  "0 : Fin(A)"
    consI   "[| a: A;  b: Fin(A) |] ==> cons(a,b) : Fin(A)"
  type_intrs  empty_subsetI, cons_subsetI, PowI
  type_elims "[make_elim PowD]"
end
```

Theory `Finite` extends the parent theory `Arith` by declaring the unary func-
tion symbol `Fin`, which is defined inductively. Its domain is specified as
$\mathcal{P}(A)$, where A is the parameter appearing in the introduction rules. For
type-checking, we supply two introduction rules:

$$\emptyset \subseteq A \qquad \frac{a \in C \quad B \subseteq C}{\{a\} \cup B \subseteq C}$$

A further introduction rule and an elimination rule express both directions
of the equivalence $A \in \mathcal{P}(B) \leftrightarrow A \subseteq B$. Type-checking involves mostly
introduction rules.

Like all Isabelle theory files, this one yields a structure containing the new
theory as an ML value. Structure `Finite` also has a substructure, called `Fin`.
After declaring `open Finite;` we can refer to the $\text{Fin}(A)$ introduction rules
as the list `Fin.intrs` or individually as `Fin.emptyI` and `Fin.consI`. The
induction rule is `Fin.induct`.

5.2   Lists of n elements

This has become a standard example of an inductive definition. Following
Paulin-Mohring [21], we could attempt to define a new datatype $\text{listn}(A, n)$,
for lists of length n, as an n-indexed family of sets. But her introduction
rules

$$\text{Niln} \in \text{listn}(A, 0) \qquad \frac{n \in \text{nat} \quad a \in A \quad l \in \text{listn}(A, n)}{\text{Consn}(n, a, l) \in \text{listn}(A, \text{succ}(n))}$$

are not acceptable to the inductive definition package: `listn` occurs with
three different parameter lists in the definition.

The Isabelle version of this example suggests a general treatment of vary-
ing parameters. It uses the existing datatype definition of $\text{list}(A)$, with
constructors `Nil` and `Cons`, and incorporates the parameter n into the in-
ductive set itself. It defines $\text{listn}(A)$ as a relation consisting of pairs $\langle n, l \rangle$

such that $n \in \mathtt{nat}$ and $l \in \mathtt{list}(A)$ and l has length n. In fact, $\mathtt{listn}(A)$ is the converse of the length function on $\mathtt{list}(A)$. The Isabelle/ZF introduction rules are

$$\langle 0, \mathtt{Nil} \rangle \in \mathtt{listn}(A) \qquad \frac{a \in A \quad \langle n, l \rangle \in \mathtt{listn}(A)}{\langle \mathtt{succ}(n), \mathtt{Cons}(a, l) \rangle \in \mathtt{listn}(A)}$$

The Isabelle theory file takes, as parent, the theory $\mathtt{List}$ of lists. We declare the constant $\mathtt{listn}$ and supply an inductive definition, specifying the domain as $\mathtt{nat} \times \mathtt{list}(A)$:

```
ListN = List +
consts  listn :: i=>i
inductive
  domains   "listn(A)" <= "nat*list(A)"
  intrs
    NilI  "<0,Nil>: listn(A)"
    ConsI "[| a:A; <n,l>:listn(A) |] ==> <succ(n), Cons(a,l)>: listn(A)"
  type_intrs "nat_typechecks @ list.intrs"
end
```

The type-checking rules include those for 0, $\mathtt{succ}$, $\mathtt{Nil}$ and $\mathtt{Cons}$. Because $\mathtt{listn}(A)$ is a set of pairs, type-checking requires the equivalence $\langle a, b \rangle \in A \times B \leftrightarrow a \in A \wedge b \in B$. The package always includes the rules for ordered pairs.

The package returns introduction, elimination and induction rules for $\mathtt{listn}$. The basic induction rule, $\mathtt{listn.induct}$, is

$$\frac{\langle z_1, z_2 \rangle \in \mathtt{listn}(A) \quad P(0, \mathtt{Nil}) \qquad \begin{array}{c} [a \in A \quad \langle n, l \rangle \in \mathtt{listn}(A) \quad P(n, l)]_{a,l,n} \\ \vdots \\ P(\mathtt{succ}(n), \mathtt{Cons}(a, l)) \end{array}}{P(z_1, z_2)}$$

This rule lets the induction formula to be a binary property of pairs, $P(n, l)$. It is now a simple matter to prove theorems about $\mathtt{listn}(A)$, such as

$$\forall l \in \mathtt{list}(A) \,.\, \langle \mathtt{length}(l), \, l \rangle \in \mathtt{listn}(A)$$

$$\mathtt{listn}(A)\text{``}\{n\} = \{l \in \mathtt{list}(A) \,.\, \mathtt{length}(l) = n\}$$

This latter result — here $r\text{``}X$ denotes the image of X under r — asserts that the inductive definition agrees with the obvious notion of n-element list.

A "list of n elements" really is a list, namely an element of $\mathtt{list}(A)$. It is subject to list operators such as append (concatenation). For example, a trivial induction on $\langle m, l \rangle \in \mathtt{listn}(A)$ yields

$$\frac{\langle m, l \rangle \in \mathtt{listn}(A) \quad \langle m', l' \rangle \in \mathtt{listn}(A)}{\langle m + m', \, l@l' \rangle \in \mathtt{listn}(A)}$$

where $+$ denotes addition on the natural numbers and $@$ denotes append.

11

5.3 Rule inversion: the function `mk_cases`

The elimination rule, `listn.elim`, is cumbersome:

$$
\cfrac{x \in \mathtt{listn}(A) \qquad \cfrac{[x = \langle 0, \mathtt{Nil} \rangle]}{\vdots \\ Q} \qquad \cfrac{\left[\begin{array}{l} x = \langle \mathtt{succ}(n), \mathtt{Cons}(a, l) \rangle \\ a \in A \\ \langle n, l \rangle \in \mathtt{listn}(A) \end{array}\right]_{a,l,n}}{\vdots \\ Q}}{Q}
$$

The ML function `listn.mk_cases` generates simplified instances of this rule. It works by freeness reasoning on the list constructors: $\mathtt{Cons}(a, l)$ is injective in its two arguments and differs from `Nil`. If x is $\langle i, \mathtt{Nil} \rangle$ or $\langle i, \mathtt{Cons}(a, l) \rangle$ then `listn.mk_cases` deduces the corresponding form of i; this is called rule inversion. Here is a sample session:

```
listn.mk_cases "<i,Nil> : listn(A)";
  "[| <?i, []> : listn(?A); ?i = 0 ==> ?Q |] ==> ?Q" : thm
listn.mk_cases "<i,Cons(a,l)> : listn(A)";
  "[| <?i, Cons(?a, ?l)> : listn(?A);
      !!n. [| ?a : ?A; <n, ?l> : listn(?A); ?i = succ(n) |] ==> ?Q
  |] ==> ?Q" : thm
```

Each of these rules has only two premises. In conventional notation, the second rule is

$$
\cfrac{\langle i, \mathtt{Cons}(a, l) \rangle \in \mathtt{listn}(A) \qquad \cfrac{\left[\begin{array}{l} a \in A \\ \langle n, l \rangle \in \mathtt{listn}(A) \\ i = \mathtt{succ}(n) \end{array}\right]_{n}}{\vdots \\ Q}}{Q}
$$

The package also has built-in rules for freeness reasoning about 0 and `succ`. So if x is $\langle 0, l \rangle$ or $\langle \mathtt{succ}(i), l \rangle$, then `listn.mk_cases` can deduce the corresponding form of l.

The function `mk_cases` is also useful with datatype definitions. The instance from the definition of lists, namely `list.mk_cases`, can prove that $\mathtt{Cons}(a, l) \in \mathtt{list}(A)$ implies $a \in A$ and $l \in \mathtt{list}(A)$:

$$
\cfrac{\mathtt{Cons}(a, l) \in \mathtt{list}(A) \qquad \cfrac{[a \in A \quad l \in \mathtt{list}(A)]}{\vdots \\ Q}}{Q}
$$

A typical use of `mk_cases` concerns inductive definitions of evaluation relations. Then rule inversion yields case analysis on possible evaluations. For example, Isabelle/ZF includes a short proof of the diamond property for parallel contraction on combinators. Ole Rasmussen used `mk_cases` extensively in his development of the theory of residuals [31].

5.4   A coinductive definition: bisimulations on lazy lists

This example anticipates the definition of the codatatype `llist`(A), which consists of finite and infinite lists over A. Its constructors are `LNil` and `LCons`, satisfying the introduction rules shown in §4.3. Because `llist`(A) is defined as a greatest fixedpoint and uses the variant pairing and injection operators, it contains non-well-founded elements such as solutions to `LCons`$(a, l) = l$.

The next step in the development of lazy lists is to define a coinduction principle for proving equalities. This is done by showing that the equality relation on lazy lists is the greatest fixedpoint of some monotonic operation. The usual approach [30] is to define some notion of bisimulation for lazy lists, define equivalence to be the greatest bisimulation, and finally to prove that two lazy lists are equivalent if and only if they are equal. The coinduction rule for equivalence then yields a coinduction principle for equalities.

A binary relation R on lazy lists is a **bisimulation** provided $R \subseteq R^+$, where R^+ is the relation

$$\{\langle \texttt{LNil}, \texttt{LNil} \rangle\} \cup \{\langle \texttt{LCons}(a, l), \texttt{LCons}(a, l') \rangle . a \in A \wedge \langle l, l' \rangle \in R\}.$$

A pair of lazy lists are **equivalent** if they belong to some bisimulation. Equivalence can be coinductively defined as the greatest fixedpoint for the introduction rules

$$\langle \texttt{LNil}, \texttt{LNil} \rangle \in \texttt{lleq}(A) \qquad \frac{a \in A \quad \langle l, l' \rangle \in \texttt{lleq}(A)}{\langle \texttt{LCons}(a, l), \texttt{LCons}(a, l') \rangle \in \texttt{lleq}(A)} \; (-)$$

To make this coinductive definition, the theory file includes (after the declaration of `llist`(A)) the following lines:

```
consts    lleq :: i=>i
coinductive
  domains "lleq(A)" <= "llist(A) * llist(A)"
  intrs
    LNil  "<LNil,LNil> : lleq(A)"
    LCons "[| a:A; <l,l'>:lleq(A) |] ==> <LCons(a,l),LCons(a,l')>: lleq(A)"
  type_intrs  "llist.intrs"
```

The domain of `lleq`(A) is `llist`$(A) \times$ `llist`(A). The type-checking rules include the introduction rules for `llist`(A), whose declaration is discussed below (§6.3).

The package returns the introduction rules and the elimination rule, as usual. But instead of induction rules, it returns a coinduction rule. The rule is too big to display in the usual notation; its conclusion is $x \in \mathtt{lleq}(A)$ and its premises are $x \in X$, $X \subseteq \mathtt{llist}(A) \times \mathtt{llist}(A)$ and

$$
\begin{array}{c}
[z \in X]_z \\
\vdots \\
z = \langle \mathtt{LNil}, \mathtt{LNil} \rangle \vee (\exists a\, l\, l' \,.\, z = \langle \mathtt{LCons}(a,l), \mathtt{LCons}(a,l') \rangle \wedge a \in A \wedge \\
\langle l, l' \rangle \in X \cup \mathtt{lleq}(A))
\end{array}
$$

Thus if $x \in X$, where X is a bisimulation contained in the domain of $\mathtt{lleq}(A)$, then $x \in \mathtt{lleq}(A)$. It is easy to show that $\mathtt{lleq}(A)$ is reflexive: the equality relation is a bisimulation. And $\mathtt{lleq}(A)$ is symmetric: its converse is a bisimulation. But showing that $\mathtt{lleq}(A)$ coincides with the equality relation takes some work.

5.5 The accessible part of a relation

Let $\prec$ be a binary relation on D; in short, $(\prec) \subseteq D \times D$. The **accessible** or **well-founded** part of $\prec$, written $\mathtt{acc}(\prec)$, is essentially that subset of D for which $\prec$ admits no infinite decreasing chains [2]. Formally, $\mathtt{acc}(\prec)$ is inductively defined to be the least set that contains a if it contains all $\prec$-predecessors of a, for $a \in D$. Thus we need an introduction rule of the form

$$
\frac{\forall y \,.\, y \prec a \rightarrow y \in \mathtt{acc}(\prec)}{a \in \mathtt{acc}(\prec)}
$$

Paulin-Mohring treats this example in Coq [21], but it causes difficulties for other systems. Its premise is not acceptable to the inductive definition package of the Cambridge HOL system [5]. It is also unacceptable to the Isabelle package (recall §3.1), but fortunately can be transformed into the acceptable form $t \in M(R)$.

The powerset operator is monotonic, and $t \in \mathcal{P}(R)$ is equivalent to $t \subseteq R$. This in turn is equivalent to $\forall y \in t \,.\, y \in R$. To express $\forall y \,.\, y \prec a \rightarrow y \in \mathtt{acc}(\prec)$ we need only find a term t such that $y \in t$ if and only if $y \prec a$. A suitable t is the inverse image of $\{a\}$ under $\prec$.

The definition below follows this approach. Here r is $\prec$ and $\mathtt{field}(r)$ refers to D, the domain of $\mathtt{acc}(r)$. (The field of a relation is the union of its domain and range.) Finally $r^-\text{``}\{a\}$ denotes the inverse image of $\{a\}$ under r. We supply the theorem $\mathtt{Pow\_mono}$, which asserts that $\mathcal{P}$ is monotonic.

```
consts     acc :: i=>i
inductive
  domains "acc(r)" <= "field(r)"
  intrs
    vimage  "[| r-''{a}: Pow(acc(r)); a: field(r) |] ==> a: acc(r)"
  monos       Pow_mono
```

The Isabelle theory proceeds to prove facts about $\mathtt{acc}(\prec)$. For instance, $\prec$ is well-founded if and only if its field is contained in $\mathtt{acc}(\prec)$.

As mentioned in §4.1, a premise of the form $t \in M(R)$ gives rise to an unusual induction hypothesis. Let us examine the induction rule, $\mathtt{acc.induct}$:

$$\frac{x \in \mathtt{acc}(r) \qquad \overset{\displaystyle \left[\begin{array}{c} r^{-\,\text{\textquotedblleft}}\{a\} \in \mathcal{P}(\{z \in \mathtt{acc}(r)\,.\,P(z)\}) \\ a \in \mathtt{field}(r) \end{array}\right]_a}{\vdots} \atop P(a)}{P(x)}$$

The strange induction hypothesis is equivalent to $\forall y\,.\,\langle y, a\rangle \in r \to y \in \mathtt{acc}(r) \wedge P(y)$. Therefore the rule expresses well-founded induction on the accessible part of $\prec$.

The use of inverse image is not essential. The Isabelle package can accept introduction rules with arbitrary premises of the form $\forall \vec{y}\,.\,P(\vec{y}) \to f(\vec{y}) \in R$. The premise can be expressed equivalently as

$$\{z \in D\,.\,P(\vec{y}) \wedge z = f(\vec{y})\} \in \mathcal{P}(R)$$

provided $f(\vec{y}) \in D$ for all $\vec{y}$ such that $P(\vec{y})$. The following section demonstrates another use of the premise $t \in M(R)$, where $M = \mathtt{list}$.

5.6   The primitive recursive functions

The primitive recursive functions are traditionally defined inductively, as a subset of the functions over the natural numbers. One difficulty is that functions of all arities are taken together, but this is easily circumvented by regarding them as functions on lists. Another difficulty, the notion of composition, is less easily circumvented.

Here is a more precise definition. Letting $\vec{x}$ abbreviate $x_0, \ldots, x_{n-1}$, we can write lists such as $[\vec{x}]$, $[y+1, \vec{x}]$, etc. A function is **primitive recursive** if it belongs to the least set of functions in $\mathtt{list(nat)} \to \mathtt{nat}$ containing

- The **successor** function $\mathtt{SC}$, such that $\mathtt{SC}[y, \vec{x}] = y + 1$.

- All **constant** functions $\mathtt{CONST}(k)$, such that $\mathtt{CONST}(k)[\vec{x}] = k$.

15

- All **projection** functions $\texttt{PROJ}(i)$, such that $\texttt{PROJ}(i)[\vec{x}] = x_i$ if $0 \le i < n$.

- All **compositions** $\texttt{COMP}(g, [f_0, \ldots, f_{m-1}])$, where g and $f_0, \ldots, f_{m-1}$ are primitive recursive, such that

$$\texttt{COMP}(g, [f_0, \ldots, f_{m-1}])[\vec{x}] = g[f_0[\vec{x}], \ldots, f_{m-1}[\vec{x}]].$$

- All **recursions** $\texttt{PREC}(f, g)$, where f and g are primitive recursive, such that

$$
\begin{aligned}
\texttt{PREC}(f, g)[0, \vec{x}] &= f[\vec{x}] \\
\texttt{PREC}(f, g)[y + 1, \vec{x}] &= g[\texttt{PREC}(f, g)[y, \vec{x}],\ y,\ \vec{x}].
\end{aligned}
$$

Composition is awkward because it combines not two functions, as is usual, but $m+1$ functions. In her proof that Ackermann's function is not primitive recursive, Nora Szasz was unable to formalize this definition directly [34]. So she generalized primitive recursion to tuple-valued functions. This modified the inductive definition such that each operation on primitive recursive functions combined just two functions.

Szasz was using ALF, but Coq and HOL would also have problems accepting this definition. Isabelle's package accepts it easily since $[f_0, \ldots, f_{m-1}]$ is a list of primitive recursive functions and `list` is monotonic. There are five introduction rules, one for each of the five forms of primitive recursive function. Let us examine the one for `COMP`:

$$\frac{g \in \texttt{primrec} \quad fs \in \texttt{list}(\texttt{primrec})}{\texttt{COMP}(g, fs) \in \texttt{primrec}}$$

The induction rule for `primrec` has one case for each introduction rule. Due to the use of `list` as a monotone operator, the composition case has an unusual induction hypothesis:

$$
\begin{array}{c}
[g \in \texttt{primrec} \quad fs \in \texttt{list}(\{z \in \texttt{primrec} \,.\, P(z)\})]_{fs,g} \\
\vdots \\
P(\texttt{COMP}(g, fs))
\end{array}
$$

The hypothesis states that fs is a list of primitive recursive functions, each satisfying the induction formula. Proving the `COMP` case typically requires structural induction on lists, yielding two subcases: either $fs = \texttt{Nil}$ or else $fs = \texttt{Cons}(f, fs')$, where $f \in \texttt{primrec}$, $P(f)$, and fs' is another list of primitive recursive functions satisfying P.

```
Primrec_defs = Main +
consts SC :: i
 .
 .
 .
defs
 SC_def     "SC == lam l:list(nat).list_case(0, %x xs.succ(x), l)"
 .
 .
 .
end

Primrec = Primrec_defs +
consts prim_rec :: i
inductive
 domains "primrec" <= "list(nat)->nat"
 intrs
   SC      "SC : primrec"
   CONST   "k: nat ==> CONST(k) : primrec"
   PROJ    "i: nat ==> PROJ(i) : primrec"
   COMP    "[| g: primrec; fs: list(primrec) |] ==> COMP(g,fs): primrec"
   PREC    "[| f: primrec; g: primrec |] ==> PREC(f,g): primrec"
 monos       list_mono
 con_defs    SC_def, CONST_def, PROJ_def, COMP_def, PREC_def
 type_intrs "nat_typechecks @ list.intrs @
             [lam_type, list_case_type, drop_type, map_type,
              apply_type, rec_type]"
end
```

Figure 1: Inductive definition of the primitive recursive functions

Figure 1 presents the theory file. Theory `Primrec` defines the constants `SC`, `CONST`, etc. These are not constructors of a new datatype, but functions over lists of numbers. Their definitions, most of which are omitted, consist of routine list programming. In Isabelle/ZF, the primitive recursive functions are defined as a subset of the function set $\texttt{list(nat)} \to \texttt{nat}$.

The Isabelle theory goes on to formalize Ackermann's function and prove that it is not primitive recursive, using the induction rule `primrec.induct`. The proof follows Szasz's excellent account.

6 Datatypes and codatatypes

A (co)datatype definition is a (co)inductive definition with automatically defined constructors and a case analysis operator. The package proves that the case operator inverts the constructors and can prove freeness theorems involving any pair of constructors.

6.1 Constructors and their domain

A (co)inductive definition selects a subset of an existing set; a (co)datatype definition creates a new set. The package reduces the latter to the former. Isabelle/ZF supplies sets having strong closure properties to serve as domains for (co)inductive definitions.

Isabelle/ZF defines the Cartesian product $A \times B$, containing ordered pairs $\langle a, b \rangle$; it also defines the disjoint sum $A + B$, containing injections $\texttt{Inl}(a) \equiv \langle 0, a \rangle$ and $\texttt{Inr}(b) \equiv \langle 1, b \rangle$. For use below, define the m-tuple $\langle x_1, \ldots, x_m \rangle$ to be the empty set $\emptyset$ if $m = 0$, simply x_1 if $m = 1$ and $\langle x_1, \langle x_2, \ldots, x_m \rangle \rangle$ if $m \geq 2$.

A datatype constructor $\texttt{Con}(x_1, \ldots, x_m)$ is defined to be $h(\langle x_1, \ldots, x_m \rangle)$, where h is composed of `Inl` and `Inr`. In a mutually recursive definition, all constructors for the set R_i have the outer form h_{in}, where h_{in} is the injection described in §3.3. Further nested injections ensure that the constructors for R_i are pairwise distinct.

Isabelle/ZF defines the set $\texttt{univ}(A)$, which contains A and furthermore contains $\langle a, b \rangle$, $\texttt{Inl}(a)$ and $\texttt{Inr}(b)$ for $a, b \in \texttt{univ}(A)$. In a typical datatype definition with set parameters A_1, ..., A_k, a suitable domain for all the recursive sets is $\texttt{univ}(A_1 \cup \cdots \cup A_k)$. This solves the problem for datatypes [25, §4.2].

The standard pairs and injections can only yield well-founded constructions. This eases the (manual!) definition of recursive functions over data-

types. But they are unsuitable for codatatypes, which typically contain non-well-founded objects.

To support codatatypes, Isabelle/ZF defines a variant notion of ordered pair, written $\langle a; b \rangle$. It also defines the corresponding variant notion of Cartesian product $A \otimes B$, variant injections $\mathtt{QInl}(a)$ and $\mathtt{QInr}(b)$ and variant disjoint sum $A \oplus B$. Finally it defines the set $\mathtt{quniv}(A)$, which contains A and furthermore contains $\langle a; b \rangle$, $\mathtt{QInl}(a)$ and $\mathtt{QInr}(b)$ for a, $b \in \mathtt{quniv}(A)$. In a typical codatatype definition with set parameters A_1, ..., A_k, a suitable domain is $\mathtt{quniv}(A_1 \cup \cdots \cup A_k)$. Details are published elsewhere [27].

6.2 The case analysis operator

The (co)datatype package automatically defines a case analysis operator, called $R\_\mathtt{case}$. A mutually recursive definition still has only one operator, whose name combines those of the recursive sets: it is called $R_1\_\ldots\_R_n\_\mathtt{case}$. The case operator is analogous to those for products and sums.

Datatype definitions employ standard products and sums, whose operators are $\mathtt{split}$ and $\mathtt{case}$ and satisfy the equations

$$
\begin{aligned}
\mathtt{split}(f, \langle x, y \rangle) &= f(x, y) \\
\mathtt{case}(f, g, \mathtt{Inl}(x)) &= f(x) \\
\mathtt{case}(f, g, \mathtt{Inr}(y)) &= g(y)
\end{aligned}
$$

Suppose the datatype has k constructors $\mathtt{Con}_1$, ..., $\mathtt{Con}_k$. Then its case operator takes $k+1$ arguments and satisfies an equation for each constructor:

$$
R\_\mathtt{case}(f_1, \ldots, f_k, \mathtt{Con}_i(\vec{x})) = f_i(\vec{x}), \qquad i = 1, \ldots, k
$$

The case operator's definition takes advantage of Isabelle's representation of syntax in the typed λ-calculus; it could readily be adapted to a theorem prover for higher-order logic. If f and g have meta-type $i \Rightarrow i$ then so do $\mathtt{split}(f)$ and $\mathtt{case}(f, g)$. This works because $\mathtt{split}$ and $\mathtt{case}$ operate on their last argument. They are easily combined to make complex case analysis operators. For example, $\mathtt{case}(f, \mathtt{case}(g, h))$ performs case analysis for $A + (B + C)$; let us verify one of the three equations:

$$
\mathtt{case}(f, \mathtt{case}(g, h), \mathtt{Inr}(\mathtt{Inl}(b))) = \mathtt{case}(g, h, \mathtt{Inl}(b)) = g(b)
$$

Codatatype definitions are treated in precisely the same way. They express case operators using those for the variant products and sums, namely $\mathtt{qsplit}$ and $\mathtt{qcase}$.

To see how constructors and the case analysis operator are defined, let us examine some examples. Further details are available elsewhere [25].

19

6.3 Example: lists and lazy lists

Here is a declaration of the datatype of lists, as it might appear in a theory file:

```
consts  list :: i=>i
datatype "list(A)" = Nil | Cons ("a:A", "l: list(A)")
```

And here is a declaration of the codatatype of lazy lists:

```
consts  llist :: i=>i
codatatype "llist(A)" = LNil | LCons ("a: A", "l: llist(A)")
```

Each form of list has two constructors, one for the empty list and one for adding an element to a list. Each takes a parameter, defining the set of lists over a given set A. Each is automatically given the appropriate domain: $\mathtt{univ}(A)$ for $\mathtt{list}(A)$ and $\mathtt{quniv}(A)$ for $\mathtt{llist}(A)$. The default can be overridden.

Since $\mathtt{list}(A)$ is a datatype, it has a structural induction rule, $\mathtt{list.induct}$:

$$\frac{x \in \mathtt{list}(A) \quad P(\mathtt{Nil}) \qquad \begin{array}{c} [a \in A \quad l \in \mathtt{list}(A) \quad P(l)]_{a,l} \\ \vdots \\ P(\mathtt{Cons}(a,l)) \end{array}}{P(x)}$$

Induction and freeness yield the law $l \neq \mathtt{Cons}(a,l)$. To strengthen this, Isabelle/ZF defines the rank of a set and proves that the standard pairs and injections have greater rank than their components. An immediate consequence, which justifies structural recursion on lists [25, §4.3], is

$$\mathtt{rank}(l) < \mathtt{rank}(\mathtt{Cons}(a,l)).$$

But $\mathtt{llist}(A)$ is a codatatype and has no induction rule. Instead it has the coinduction rule shown in §4.3. Since variant pairs and injections are monotonic and need not have greater rank than their components, fixedpoint operators can create cyclic constructions. For example, the definition

$$\mathtt{lconst}(a) \equiv \mathtt{lfp}(\mathtt{univ}(a), \lambda l \,.\, \mathtt{LCons}(a,l))$$

yields $\mathtt{lconst}(a) = \mathtt{LCons}(a, \mathtt{lconst}(a))$.

It may be instructive to examine the definitions of the constructors and case operator for $\mathtt{list}(A)$. The definitions for $\mathtt{llist}(A)$ are similar. The list constructors are defined as follows:

$$\begin{aligned} \mathtt{Nil} &\equiv \mathtt{Inl}(\emptyset) \\ \mathtt{Cons}(a,l) &\equiv \mathtt{Inr}(\langle a, l \rangle) \end{aligned}$$

The operator `list_case` performs case analysis on these two alternatives:

$$\texttt{list\_case}(c, h) \equiv \texttt{case}(\lambda u \mathbin{.} c, \texttt{split}(h))$$

Let us verify the two equations:

$$
\begin{aligned}
\texttt{list\_case}(c, h, \texttt{Nil}) &= \texttt{case}(\lambda u \mathbin{.} c, \texttt{split}(h), \texttt{Inl}(\emptyset)) \\
&= (\lambda u \mathbin{.} c)(\emptyset) \\
&= c \\
\texttt{list\_case}(c, h, \texttt{Cons}(x, y)) &= \texttt{case}(\lambda u \mathbin{.} c, \texttt{split}(h), \texttt{Inr}(\langle x, y \rangle)) \\
&= \texttt{split}(h, \langle x, y \rangle) \\
&= h(x, y)
\end{aligned}
$$

6.4   Example: mutual recursion

In mutually recursive trees and forests [25, §4.5], trees have the one constructor `Tcons`, while forests have the two constructors `Fnil` and `Fcons`:

```
consts  tree, forest, tree_forest    :: i=>i
datatype "tree(A)"   = Tcons ("a: A",  "f: forest(A)")
and      "forest(A)" = Fnil  |  Fcons ("t: tree(A)",  "f: forest(A)")
```

The three introduction rules define the mutual recursion. The distinguishing feature of this example is its two induction rules.

The basic induction rule is called `tree_forest.induct`:

$$
\cfrac{
x \in \texttt{tree\_forest}(A) \quad
\begin{bmatrix} a \in A \\ f \in \texttt{forest}(A) \\ P(f) \end{bmatrix}_{a,f} \\ \vdots \\ P(\texttt{Tcons}(a, f)) \quad
P(\texttt{Fnil}) \quad
\begin{bmatrix} t \in \texttt{tree}(A) \\ P(t) \\ f \in \texttt{forest}(A) \\ P(f) \end{bmatrix}_{t,f} \\ \vdots \\ P(\texttt{Fcons}(t, f))
}{P(x)}
$$

This rule establishes a single predicate for `tree_forest`(A), the union of the recursive sets. Although such reasoning can be useful [25, §4.5], a proper mutual induction rule should establish separate predicates for `tree`(A) and `forest`(A). The package calls this rule `tree_forest.mutual_induct`. Ob-

serve the usage of P and Q in the induction hypotheses:

$$\frac{\begin{bmatrix} a \in A \\ f \in \texttt{forest}(A) \\ Q(f) \end{bmatrix}_{a,f} \quad\quad \begin{bmatrix} t \in \texttt{tree}(A) \\ P(t) \\ f \in \texttt{forest}(A) \\ Q(f) \end{bmatrix}_{t,f} }{ \begin{array}{ccc} \vdots & & \vdots \\ P(\texttt{Tcons}(a,f)) & Q(\texttt{Fnil}) & Q(\texttt{Fcons}(t,f)) \end{array} }$$
$$\overline{(\forall z \,.\, z \in \texttt{tree}(A) \rightarrow P(z)) \wedge (\forall z \,.\, z \in \texttt{forest}(A) \rightarrow Q(z))}$$

Elsewhere I describe how to define mutually recursive functions over trees and forests [25, §4.5].

Both forest constructors have the form $\texttt{Inr}(\cdots)$, while the tree constructor has the form $\texttt{Inl}(\cdots)$. This pattern would hold regardless of how many tree or forest constructors there were.

$$
\begin{aligned}
\texttt{Tcons}(a,l) &\equiv \texttt{Inl}(\langle a,l \rangle) \\
\texttt{Fnil} &\equiv \texttt{Inr}(\texttt{Inl}(\emptyset)) \\
\texttt{Fcons}(a,l) &\equiv \texttt{Inr}(\texttt{Inr}(\langle a,l \rangle))
\end{aligned}
$$

There is only one case operator; it works on the union of the trees and forests:

$$\texttt{tree\_forest\_case}(f,c,g) \equiv \texttt{case}(\texttt{split}(f), \texttt{case}(\lambda u \,.\, c, \texttt{split}(g)))$$

6.5 Example: a four-constructor datatype

A bigger datatype will illustrate some efficiency refinements. It has four constructors $\texttt{Con}_0, \ldots, \texttt{Con}_3$, with the corresponding arities.

```
consts    data :: [i,i] => i
datatype  "data(A,B)" = Con0
                      | Con1 ("a: A")
                      | Con2 ("a: A", "b: B")
                      | Con3 ("a: A", "b: B", "d: data(A,B)")
```

Because this datatype has two set parameters, A and B, the package automatically supplies $\texttt{univ}(A \cup B)$ as its domain. The structural induction rule has four minor premises, one per constructor, and only the last has an induction hypothesis. (Details are left to the reader.)

The constructors are defined by the equations

$$
\begin{aligned}
\texttt{Con}_0 &\equiv \texttt{Inl}(\texttt{Inl}(\emptyset)) \\
\texttt{Con}_1(a) &\equiv \texttt{Inl}(\texttt{Inr}(a)) \\
\texttt{Con}_2(a,b) &\equiv \texttt{Inr}(\texttt{Inl}(\langle a,b \rangle)) \\
\texttt{Con}_3(a,b,c) &\equiv \texttt{Inr}(\texttt{Inr}(\langle a,b,c \rangle)).
\end{aligned}
$$

The case analysis operator is

$$\texttt{data\_case}(f_0, f_1, f_2, f_3) \equiv \texttt{case}(\texttt{case}(\lambda u \,.\, f_0, \ f_1),$$
$$\texttt{case}(\texttt{split}(f_2), \ \texttt{split}(\lambda v \,.\, \texttt{split}(f_3(v)))))$$

This may look cryptic, but the case equations are trivial to verify.

In the constructor definitions, the injections are balanced. A more naive approach is to define $\texttt{Con}_3(a, b, c)$ as $\texttt{Inr}(\texttt{Inr}(\texttt{Inr}(\langle a, b, c \rangle)))$; instead, each constructor has two injections. The difference here is small. But the ZF examples include a 60-element enumeration type, where each constructor has 5 or 6 injections. The naive approach would require 1 to 59 injections; the definitions would be quadratic in size. It is like the advantage of binary notation over unary.

The result structure contains the case operator and constructor definitions as the theorem list `con_defs`. It contains the case equations, such as

$$\texttt{data\_case}(f_0, f_1, f_2, f_3, \texttt{Con}_3(a, b, c)) = f_3(a, b, c),$$

as the theorem list `case_eqns`. There is one equation per constructor.

6.6 Proving freeness theorems

There are two kinds of freeness theorems:

- **injectiveness** theorems, such as

$$\texttt{Con}_2(a, b) = \texttt{Con}_2(a', b') \leftrightarrow a = a' \wedge b = b'$$

- **distinctness** theorems, such as

$$\texttt{Con}_1(a) \neq \texttt{Con}_2(a', b')$$

Since the number of such theorems is quadratic in the number of constructors, the package does not attempt to prove them all. Instead it returns tools for proving desired theorems — either manually or during simplification or classical reasoning.

The theorem list `free_iffs` enables the simplifier to perform freeness reasoning. This works by incremental unfolding of constructors that appear in equations. The theorem list contains logical equivalences such as

$$\texttt{Con}_0 = c \ \leftrightarrow \ c = \texttt{Inl}(\texttt{Inl}(\emptyset))$$
$$\texttt{Con}_1(a) = c \ \leftrightarrow \ c = \texttt{Inl}(\texttt{Inr}(a))$$

$$\vdots$$
$$
\begin{aligned}
\mathtt{Inl}(a) = \mathtt{Inl}(b) &\quad\leftrightarrow\quad a = b \\
\mathtt{Inl}(a) = \mathtt{Inr}(b) &\quad\leftrightarrow\quad \mathtt{False} \\
\langle a, b \rangle = \langle a', b' \rangle &\quad\leftrightarrow\quad a = a' \wedge b = b'
\end{aligned}
$$

For example, these rewrite $\mathtt{Con_1}(a) = \mathtt{Con_1}(b)$ to $a = b$ in four steps.

The theorem list `free_SEs` enables the classical reasoner to perform similar replacements. It consists of elimination rules to replace $\mathtt{Con_0} = c$ by $c = \mathtt{Inl(Inl(\emptyset))}$ and so forth, in the assumptions.

Such incremental unfolding combines freeness reasoning with other proof steps. It has the unfortunate side-effect of unfolding definitions of constructors in contexts such as $\exists x \,.\, \mathtt{Con_1}(a) = x$, where they should be left alone. Calling the Isabelle tactic `fold_tac con_defs` restores the defined constants.

7 Related work

The use of least fixedpoints to express inductive definitions seems obvious. Why, then, has this technique so seldom been implemented?

Most automated logics can only express inductive definitions by asserting axioms. Little would be left of Boyer and Moore's logic [4] if their shell principle were removed. With ALF the situation is more complex; earlier versions of Martin-Löf's type theory could (using wellordering types) express datatype definitions, but the version underlying ALF requires new rules for each definition [7]. With Coq the situation is subtler still; its underlying Calculus of Constructions can express inductive definitions [13], but cannot quite handle datatype definitions [21]. It seems that researchers tried hard to circumvent these problems before finally extending the Calculus with rule schemes for strictly positive operators. Recently Giménez has extended the Calculus of Constructions with inductive and coinductive types [10], with mechanized support in Coq.

Higher-order logic can express inductive definitions through quantification over unary predicates. The following formula expresses that i belongs to the least set containing 0 and closed under `succ`:

$$\forall P \,.\, P(0) \wedge (\forall x \,.\, P(x) \rightarrow P(\mathtt{succ}(x))) \rightarrow P(i)$$

This technique can be used to prove the Knaster-Tarski theorem, which (in its general form) is little used in the Cambridge HOL system. Melham [14] describes the development. The natural numbers are defined as shown above, but lists are defined as functions over the natural numbers. Unlabelled trees

are defined using Gödel numbering; a labelled tree consists of an unlabelled tree paired with a list of labels. Melham's datatype package expresses the user's datatypes in terms of labelled trees. It has been highly successful, but a fixedpoint approach might have yielded greater power with less effort.

Elsa Gunter [11] reports an ongoing project to generalize the Cambridge HOL system with mutual recursion and infinitely-branching trees. She retains many features of Melham's approach.

Melham's inductive definition package [5] also uses quantification over predicates. But instead of formalizing the notion of monotone function, it requires definitions to consist of finitary rules, a syntactic form that excludes many monotone inductive definitions.

PVS [20] is another proof assistant based on higher-order logic. It supports both inductive definitions and datatypes, apparently by asserting axioms. Datatypes may not be iterated in general, but may use recursion over the built-in `list` type.

The earliest use of least fixedpoints is probably Robin Milner's. Brian Monahan extended this package considerably [18], as did I in unpublished work.[3] LCF is a first-order logic of domain theory; the relevant fixedpoint theorem is not Knaster-Tarski but concerns fixedpoints of continuous functions over domains. LCF is too weak to express recursive predicates. The Isabelle package might be the first to be based on the Knaster-Tarski theorem.

8 Conclusions and future work

Higher-order logic and set theory are both powerful enough to express inductive definitions. A growing number of theorem provers implement one of these [8, 32]. The easiest sort of inductive definition package to write is one that asserts new axioms, not one that makes definitions and proves theorems about them. But asserting axioms could introduce unsoundness.

The fixedpoint approach makes it fairly easy to implement a package for (co)inductive definitions that does not assert axioms. It is efficient: it processes most definitions in seconds and even a 60-constructor datatype requires only a few minutes. It is also simple: The first working version took under a week to code, consisting of under 1100 lines (35K bytes) of Standard ML.

In set theory, care is needed to ensure that the inductive definition yields a set (rather than a proper class). This problem is inherent to set theory, whether or not the Knaster-Tarski theorem is employed. We must exhibit a

---

[3]The datatype package described in my LCF book [22] does *not* make definitions, but merely asserts axioms.

bounding set (called a domain above). For inductive definitions, this is often trivial. For datatype definitions, I have had to formalize much set theory. To justify infinitely-branching datatype definitions, I have had to develop a theory of cardinal arithmetic [29], such as the theorem that if κ is an infinite cardinal and $|X(\alpha)| \leq \kappa$ for all $\alpha < \kappa$ then $|\bigcup_{\alpha<\kappa} X(\alpha)| \leq \kappa$. The need for such efforts is not a drawback of the fixedpoint approach, for the alternative is to take such definitions on faith.

Care is also needed to ensure that the greatest fixedpoint really yields a coinductive definition. In set theory, standard pairs admit only well-founded constructions. Aczel's anti-foundation axiom [3] could be used to get non-well-founded objects, but it does not seem easy to mechanize. Isabelle/ZF instead uses a variant notion of ordered pairing, which can be generalized to a variant notion of function. Elsewhere I have proved that this simple approach works (yielding final coalgebras) for a broad class of definitions [27].

Several large studies make heavy use of inductive definitions. Lötzbeyer and Sandner have formalized two chapters of a semantics book [35], proving the equivalence between the operational and denotational semantics of a simple imperative language. A single theory file contains three datatype definitions (of arithmetic expressions, boolean expressions and commands) and three inductive definitions (the corresponding operational rules). Using different techniques, Nipkow [19] and Rasmussen [31] have both proved the Church-Rosser theorem; inductive definitions specify several reduction relations on λ-terms. Recently, I have applied inductive definitions to the analysis of cryptographic protocols [28].

To demonstrate coinductive definitions, Frost [9] has proved the consistency of the dynamic and static semantics for a small functional language. The example is due to Milner and Tofte [17]. It concerns an extended correspondence relation, which is defined coinductively. A codatatype definition specifies values and value environments in mutual recursion. Non-well-founded values represent recursive functions. Value environments are variant functions from variables into values. This one key definition uses most of the package's novel features.

The approach is not restricted to set theory. It should be suitable for any logic that has some notion of set and the Knaster-Tarski theorem. I have ported the (co)inductive definition package from Isabelle/ZF to Isabelle/HOL (higher-order logic).

References

[1] Samson Abramsky. The lazy lambda calculus. In David A. Turner, editor, *Research Topics in Functional Programming*, pages 65–116. Addison-Wesley, 1990.

[2] Peter Aczel. An introduction to inductive definitions. In J. Barwise, editor, *Handbook of Mathematical Logic*, pages 739–782. North-Holland, 1977.

[3] Peter Aczel. *Non-Well-Founded Sets*. CSLI, 1988.

[4] Robert S. Boyer and J Strother Moore. *A Computational Logic*. Academic Press, 1979.

[5] J. Camilleri and T. F. Melham. Reasoning with inductively defined relations in the HOL theorem prover. Technical Report 265, Computer Laboratory, University of Cambridge, August 1992.

[6] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 1990.

[7] Peter Dybjer. Inductive sets and families in Martin-Löf's type theory and their set-theoretic semantics. In Gérard Huet and Gordon Plotkin, editors, *Logical Frameworks*, pages 280–306. Cambridge University Press, 1991.

[8] William M. Farmer, Joshua D. Guttman, and F. Javier Thayer. IMPS: An interactive mathematical proof system. *Journal of Automated Reasoning*, 11(2):213–248, 1993.

[9] Jacob Frost. A case study of co-induction in Isabelle. Technical Report 359, Computer Laboratory, University of Cambridge, February 1995.

[10] Eduardo Giménez. Codifying guarded definitions with recursive schemes. In Peter Dybjer, Bengt Nordström, and Jan Smith, editors, *Types for Proofs and Programs: International Workshop TYPES '94*, LNCS 996, pages 39–59. Springer, 1995.

[11] Elsa L. Gunter. A broader class of trees for recursive type definitions for HOL. In J. Joyce and C. Seger, editors, *Higher Order Logic Theorem Proving and Its Applications: HUG '93*, LNCS 780, pages 141–154. Springer, Published 1994.

[12] Matthew Hennessy. *The Semantics of Programming Languages: An Elementary Introduction Using Structural Operational Semantics*. Wiley, 1990.

[13] Gérard Huet. Induction principles formalized in the Calculus of Constructions. In K. Fuchi and M. Nivat, editors, *Programming of Future Generation Computers*, pages 205–216. Elsevier, 1988.

[14] Thomas F. Melham. Automating recursive type definitions in higher order logic. In Graham Birtwistle and P. A. Subrahmanyam, editors, *Current Trends in Hardware Verification and Automated Theorem Proving*, pages 341–386. Springer, 1989.

[15] Robin Milner. How to derive inductions in LCF. note, Department of Computer Science, University of Edinburgh, 1980.

[16] Robin Milner. *Communication and Concurrency*. Prentice-Hall, 1989.

[17] Robin Milner and Mads Tofte. Co-induction in relational semantics. *Theoretical Computer Science*, 87:209–220, 1991.

[18] Brian Q. Monahan. *Data Type Proofs using Edinburgh LCF*. PhD thesis, University of Edinburgh, 1984.

[19] Tobias Nipkow. More Church-Rosser proofs (in Isabelle/HOL). In M. McRobbie and J.K. Slaney, editors, *Automated Deduction — CADE-13*, volume 1104 of *Lecture Notes in Computer Science*, pages 733–747. Springer-Verlag, 1996.

[20] S. Owre, N. Shankar, and J. M. Rushby. *The PVS specification language.* Computer Science Laboratory, SRI International, Menlo Park, CA, April 1993. Beta release.

[21] Christine Paulin-Mohring. Inductive definitions in the system Coq: Rules and properties. In M. Bezem and J.F. Groote, editors, *Typed Lambda Calculi and Applications*, LNCS 664, pages 328–345. Springer, 1993.

[22] Lawrence C. Paulson. *Logic and Computation: Interactive proof with Cambridge LCF.* Cambridge University Press, 1987.

[23] Lawrence C. Paulson. Set theory for verification: I. From foundations to functions. *Journal of Automated Reasoning*, 11(3):353–389, 1993.

[24] Lawrence C. Paulson. *Isabelle: A Generic Theorem Prover.* Springer, 1994. LNCS 828.

[25] Lawrence C. Paulson. Set theory for verification: II. Induction and recursion. *Journal of Automated Reasoning*, 15(2):167–215, 1995.

[26] Lawrence C. Paulson. Mechanizing coinduction and corecursion in higher-order logic. *Journal of Logic and Computation*, 7(2):175–204, March 1997.

[27] Lawrence C. Paulson. Final coalgebras as greatest fixed points in ZF set theory. *Mathematical Structures in Computer Science*, 9(5):545–567, 1999.

[28] Lawrence C. Paulson. Tool support for logics of programs. In Manfred Broy, editor, *Mathematical Methods in Program Development: Summer School Marktoberdorf 1996*, NATO ASI Series F, pages 461–498. Springer, Published 1997.

[29] Lawrence C. Paulson and Krzysztof Grąbczewski. Mechanizing set theory: Cardinal arithmetic and the axiom of choice. *Journal of Automated Reasoning*, 17(3):291–323, December 1996.

[30] Andrew M. Pitts. A co-induction principle for recursively defined domains. *Theoretical Computer Science*, 124:195–219, 1994.

[31] Ole Rasmussen. The Church-Rosser theorem in Isabelle: A proof porting experiment. Technical Report 364, Computer Laboratory, University of Cambridge, May 1995.

[32] Mark Saaltink, Sentot Kromodimoeljo, Bill Pase, Dan Craigen, and Irwin Meisels. An EVES data abstraction example. In J. C. P. Woodcock and P. G. Larsen, editors, *FME '93: Industrial-Strength Formal Methods*, volume 670 of *Lecture Notes in Computer Science*, pages 578–596. Springer-Verlag, 1993.

[33] Konrad Slind. Function definition in higher order logic. In J. von Wright, J. Grundy, and J. Harrison, editors, *Theorem Proving in Higher Order Logics: TPHOLs '96*, volume 1125 of *Lecture Notes in Computer Science*, pages 381–397. Springer-Verlag, 1996.

[34] Nora Szasz. A machine checked proof that Ackermann's function is not primitive recursive. In Gérard Huet and Gordon Plotkin, editors, *Logical Environments*, pages 317–338. Cambridge University Press, 1993.

[35] Glynn Winskel. *The Formal Semantics of Programming Languages.* MIT Press, 1993.