

Chapter 1

Proof for Functional Programming

Simon Thompson

University of Kent at Canterbury, UK

1.1 Introduction

In this chapter we examine ways in which functional programs can be proved correct. For a number of reasons this is easier for functional than for imperative programs. In the simplest cases functional programs are equations, so the language documents itself, as it were. Beyond this we often have a higher-level expression of properties, by means of equations between functions rather than values. We can also express properties which can't simply be discussed for imperative programs, using notations for lists and other algebraic data types, for instance.

The equational model gives the spirit of the method, but it needs to be modified and strengthened in various ways in order to apply to a full functional language. The pattern of the chapter will be to give a succession of refinements of the logic as further features are added to the language. These we look at now.

The defining forms of languages are more complex than simple equations. Conditional definitions (using 'guards'), pattern matching and local definitions (in `let` and `where` clauses) each add complications, not least when put together.

Reasoning cannot be completely equational. We need to be able to reason by cases, and in general to be able to prove properties of functions defined by recursion; structural induction is the mechanism here.

With general recursion — which is a feature of all languages in the current mainstream — comes the possibility of non-termination of evaluation. In a lazy language general recursion has the more profound effect of introducing infinite and partial lists and other data structures.

In both lazy and strict languages this possibility means in turn that in interpreting the meaning of programs we are forced to introduce extra values at each type. This plainly affects the way in which the logic is expressed, and its relation to familiar properties of, say, the integers.

Because of the complications which non-termination brings, there has been recent interest in terminating languages. These we address in Section 1.9 below.

In the body of the paper we give examples of program verifications, including a compiler for arithmetic expressions, a program to re-arrange such expressions and the equivalence of two infinite lists (or streams) of factorials.

A fully-fledged language will allow users to interact with the environment in various ways, but at its simplest by reading input and writing output. This is supported in a variety of ways, including the side-effecting functions of Standard ML and the monads of Haskell 1.4. SML also allows mutable references and exceptions. In this paper we cover only the pure parts of languages, but refer readers to [2] for a perspicacious discussion of program verification for various forms of input/output including monadic IO. Recent work on modelling SML-style references can be found in [8].

1.2 The basis of functional programming: equations

In this section we examine the basis of functional programming and show how the definitions of a simple functional program can be interpreted as logical equations. An examination of how this approach can be modified and extended to work in general forms the main part of the paper.

A functional program consists of a collection of definitions of functions and other values. An example program using the notation of the Haskell language [7] is

```
test :: Integer
test = 42

id :: t -> t
id x = x

plusOne :: Integer -> Integer
plusOne n = (n+1)

minusOne :: Integer -> Integer
minusOne n = (n-1)
```

Execution of a program consists of evaluating an expression which uses the functions and other objects defined in the program (together with the built in operations of the language). Evaluation works by the replacement of sub-expressions by their values, and is complete when a value is produced. For instance, evaluation of

```
plusOne (minusOne test)
```

will proceed thus:

```
plusOne (minusOne test)
=> (minusOne test) + 1
=> (test - 1) + 1
=> (42 - 1) + 1
=> 41 + 1
=> 42
```

where it can be seen that at each stage of the evaluation one of the defining equations is used to rewrite a sub-expression which matches the left-hand side of a definition, like

```
minusOne test
```

to the corresponding right-hand side,

```
test - 1
```

The model of evaluation for a real language such as Haskell or ML is somewhat more complex; this will be reflected by the discussion in subsequent sections.

In each step of an evaluation such as this equals are replaced by equals, and this points to the basis of a logical approach to reading functional programs. The use of the equals sign in function definitions is indeed suggestive, and we can read the definitions as **logical statements** of the properties of the defined functions, thus:

$$\text{id } x \equiv x \tag{1}$$

for all x of type τ , and so on. Note that we have used the symbol ‘ \equiv ’ here for logical equality to distinguish it from both the ‘definitional’ equality used to define objects in the language, $=$, and the ‘calculational’ Boolean equality operation of the language, $==$.

Logical equations like these can be manipulated using the rules of logic in the standard way, so that we can deduce, for instance, that

$$\begin{aligned} &\text{id } (\text{id } y) \\ &\equiv \{\text{by substituting } \text{id } y \text{ for } x \text{ in (1)}\} \\ &\text{id } y \\ &\equiv \{\text{by substituting } y \text{ for } x \text{ in (1)}\} \\ &y \end{aligned}$$

In linear proofs we shall use the format above, in which the justification for each equality step of the proof is included in braces $\{\dots\}$.

So, we see a model for verification of functional programs which uses the defining equations as logical equations, and the logical laws for equality: reflexivity, symmetry, transitivity and substitution:

$$\frac{P(\mathbf{a}) \quad \mathbf{a} \equiv \mathbf{b}}{P(\mathbf{b})} (\textit{Subst})$$

to make deductions¹.

The logical versions of the definitions given here contain free variables, namely the variables of the definitions. In the remainder of the paper we will also use a closed form given by taking the universal quantification over these variables. The equation (1) will take the form $(\forall x :: \tau) (\text{id } x \equiv x)$ for example.

1.3 Pattern matching, cases and local definitions

The purely equational definition style of Section 1.2 can be made to accommodate case switches, local definitions and pattern matching by means of the appropriate higher-order combinators. Indeed, this is one way of interpreting

¹A notational aside: we use the convention that $P(\mathbf{a})$ means an expression P in which \mathbf{a} occurs; the appearance of $P(\mathbf{b})$ below the line means that the occurrences of \mathbf{a} in P have been replaced by \mathbf{b} . An alternative notation which we use for \mathbf{b} substituted for \mathbf{a} in P is $P[\mathbf{b}/\mathbf{a}]$.

the work of Bird and others, discussed further in Section 1.9.3. However, for reasons of readability and conciseness, most languages offer syntactic support for these facilities, and with this additional syntax comes the task of giving it a logical explanation.

This section gives an overview of how pattern matching, cases and local definitions are rendered logically; a more detailed examination can be found in [13], which addresses the question for Miranda. Note that here we are still considering a small (terminating) language, rather than a full language.

1.3.1 Pattern matching

Pattern matching serves to distinguish cases, as in

```
isEmptyList :: [t] -> [t]
isEmptyList [] = True
isEmptyList _ = False
```

(where ‘_’ is a wildcard pattern, matching anything), and also to allow access to the components of a compound object

```
tail :: [t] -> [t]
tail [] = []
tail (a:x) = x
```

In the example of `tail`, where the patterns do not overlap (are exclusive) and cover all eventualities (are exhaustive), the definitions can be read as logical equations.

In the general case, we need to take account of the **sequential** interpretation which is usually applied to them. Looking at `isEmptyList`, the second equation in which the ‘_’ will match any value will only be applied should the first clause not apply. We therefore need to give a description of the complement of a pattern, here `[]`, over which the remaining equations hold. The complement of `[]` will be the non-empty list, `(a:x)`, and so we can rewrite the definition of the function to give its logical form thus:

```
isEmptyList [] ≡ True
isEmptyList (a:x) ≡ False
```

As another example, consider the pattern `(a:b:x)`. This will match lists with two or more elements, and its complement is given by the two patterns `[]` and `[_]`. The full details of the way in which Miranda pattern matching definitions can be translated are to be found in [13].

1.3.2 Cases

Definitions can have alternatives depending on the (Boolean) values of guards, in a Haskell style,

```

f args
  | g1           = e1
  | g2           = e2
  | ...
  | otherwise = e

```

If the (actual values of the) parameters `args` satisfy `g1` then the result of `f args` is `e1`; should `g1` be `False` then if `g2` is `True`, `e2` is the result, and so on. In logical form we then have

$$(g_1 \equiv \text{True} \Rightarrow f \text{ args} \equiv e_1) \wedge ((g_1 \equiv \text{False} \wedge g_2 \equiv \text{True}) \Rightarrow f \text{ args} \equiv e_2) \wedge \dots$$

which renders the definition as the conjunction of a set of conditional equations.

1.3.3 Local definitions

A local definition, introduced either by a `let` or a `where` introduces a name whose scope is restricted. An example is given by the schematic

```

f :: t1 -> t2
f x = e
  where
    g :: t3 -> t4
    g y = e'

```

The function `g` is in scope in the expression `e` as well as the `where` clause. It is also important to realise that its definition will, in general, depend upon the parameter `x`. It is translated thus

$$(\forall x :: t_1) (\exists g :: t_3 \rightarrow t_4) ((\forall y :: t_3) (g y \equiv e')) \wedge f x \equiv e$$

in which the locally defined value(s) are existentially quantified, and the universal quantification over the argument values for `f` and `g` are shown explicitly. (A discussion on the form of this translation can be found in [13].)

1.3.4 Feature Interaction

The features discussed in this section can, when they appear in real programming languages such as Haskell, have complex interactions. For instance, it is not necessary to have an `otherwise` case in a guarded equation, so that it is possible for none of the guards to hold for a particular set of arguments. In this situation, the next guarded equation (and therefore pattern match) has to be examined, and this is particularly difficult to explain when the guards also refer to local definitions – a Miranda example is presented in [13].

1.3.5 Proof

The translation given here goes beyond the equational, giving axioms which involve arbitrarily deep alternations of quantifiers. In practice these quantifiers

will be stripped off, allowing conditional equational reasoning take place; the effect of the quantifications is to ensure that the scoping rules of the language are obeyed, while the conditions reflect the guards in the definitions of the language. Pattern matching is supported by the substitution mechanism of the logic.

1.4 Structural induction and recursion

In this section we consider how to strengthen our language to accommodate recursively defined functions and types while retaining the property that all computations will terminate.

At the heart of modern functional programming languages are built-in types of lists and a facility to define ‘algebraic’ data types built by the application of constructors. If we wish to build a simple-minded representation of integer arithmetic expressions — as part of a calculator or a compiler, say — we might write, using Haskell notation

```
data IntExp = Literal Int |
            Binary Op IntExp IntExp
```

```
data Op = Add | Sub | Mul
```

which describes a type whose members take two forms, built by the two constructors of the type, `Literal` and `Binary`.

- The first is `Literal n`, where `n` is an `Int` (integer).
- The second form is `Binary op ex1 ex2` where `ex1` and `ex2` are themselves `IntExps` and `op` is one of `Add`, `Sub` or `Mul` (representing three binary arithmetic operators).

An example of the type, representing the arithmetic expression $(4+3)-5$, is

```
Binary Sub (Binary Add (Literal 4) (Literal 3)) (Literal 5)
```

To define a function over `Op` it is sufficient to give its value at the three possible inputs, so that

```
opValue :: Op -> (Int -> Int -> Int)
```

```
opValue Add = (+)
opValue Sub = (-)
opValue Mul = (*)
```

serves to interpret the arithmetic operators. In a similar way, if we wish to prove that some logical property holds for all operators it is sufficient to prove that the property holds for the three values of the type.

Now, the type `IntExp` is rather more complicated, since it is recursively defined, and has an infinite number of members. However, we know that the only ways that elements are constructed are by means of a finite number of

applications of the constructors of the type. This means that an arbitrary element of the type will take one of the forms

```
Literal n
Binary op ex1 ex2
```

where `ex1` and `ex2` are themselves elements of `IntExp`.

Because every element is built up in this way, we can deduce how to define functions over `IntExp` and to prove that properties hold of all elements of `IntExp`. To define a function we use **structural recursion**, as exemplified by a function to evaluate an arithmetic expression:

```
eval :: IntExp -> Int

eval (Literal int)          = int                (2)
eval (Binary op ex1 ex2) = opValue op (eval ex1) (eval ex2)  (3)
```

Here we see the pattern of definition in which we

- give the result at `Literal int` outright; and
- give the result at `Binary op ex1 ex2` using the results already defined for `ex1` and `ex2` (as well as other components of the data value, here `op`).

It can be seen that a finite number of recursive calls will result in calls to the `Literal` case, so that functions defined in this way will be total.

In an analogous way, we can use **structural induction** to prove a property for all `IntExps`. Formally, to prove $P(e)$ for all e in `IntExp` we need to show that

Base case The property $P(\text{Literal } \text{int})$ holds for all `int`.

Induction case The property $P(\text{Binary } \text{op } \text{ex1 } \text{ex2})$ holds *on the assumption that* $P(\text{ex1})$ and $P(\text{ex2})$ hold.

Given any `IntExp t` we can see that a finite number of applications of the induction case will lead us back to the base case, and thus establish that $P(t)$ holds.

In the next section we give examples of various functions defined by structural recursion together with verification using structural induction over the `IntExp` type.

1.5 Case study: a compiler correctness proof

In this section we give a proof of correctness of a tiny compiler for arithmetic expressions using structural induction over the type of expressions, given by the algebraic data type `IntExp`. In developing the proof we explore some of the pragmatics of finding proofs.

```

data IntExp = Literal Int |
             Binary Op IntExp IntExp

data Op = Add | Sub | Mul

opValue :: Op -> (Int -> Int -> Int)

eval :: IntExp -> Int

eval (Literal int)           = int                (2)
eval (Binary op ex1 ex2) = opValue op (eval ex1) (eval ex2) (3)

```

```

data Code = PushLit Int |
           DoBinary Op

type Program = [Code]

compile :: IntExp -> Program

compile (Literal int)
    = [PushLit int]                (4)
compile (Binary op ex1 ex2)
    = compile ex1 ++ compile ex2 ++ [DoBinary op] (5)

type Stack = [Int]

run :: Program -> Stack -> Stack

run [] stack
    = stack                        (6)
run (PushLit int : program) stack
    = run program (int : stack)    (7)
run (DoBinary op : program) (v2:v1:stack)
    = run program (opValue op v1 v2 : stack) (8)

run _ _ = []                      (9)

```

Figure 1.1: A simple interpreter and compiler for expressions

Base case _____

```

run (compile (Literal int) ++ program) stack
  ≡ { by (4) }
run ([PushLit int] ++ program) stack
  ≡ { by definition of ++ }
run (PushLit int : program) stack
  ≡ { by (7) }
run program (int : stack)

run program (eval (Literal int) : stack)
  ≡ { by (2) }
run program (int : stack)

```

Induction case _____

```

run (compile (Binary op ex1 ex2) ++ program) stack
  ≡ { by (5) and associativity of ++ }
run (compile ex1 ++ compile ex2 ++ [DoBinary op] ++ program) stack
  ≡ { by the induction hypothesis for ex1 and associativity of ++ }
run (compile ex2 ++ [DoBinary op] ++ program) (eval ex1 : stack)
  ≡ { by the induction hypothesis for ex2 and associativity of ++ }
run ([DoBinary op] ++ program) (eval ex2 : eval ex1 : stack)
  ≡ { by (8) and definition of ++ }
run program (opValue op (eval ex1) (eval ex2) : stack)

run program (eval (Binary op ex1 ex2) : stack)
  ≡ { by (3) }
run program (opValue op (eval ex1) (eval ex2) : stack)

```

Correctness theorem _____

```

run (compile e) []
  ≡ { by the above }
run [] [eval e]
  ≡ { by (6) }
[eval e]

```

Figure 1.2: Proof of compiler correctness

It is instructive to compare this program and proof developed in a functional context with a similar problem programmed in a modern imperative language such as C++, Java or Modula 3. The advantage of the approach here is that modern functional languages contain explicit representations of recursive data types, and so a proof of a program property can refer explicitly to the forms of data values. In contrast, a stack in an imperative language will either be represented by a dynamic data structure, built using pointers, or by an array, with the attendant problems of working with a concrete representation of a stack rather than an appropriately abstract view. In either case it is not so easy to see how a proof could be written, indeed the most appropriate model might be to develop the imperative program by refinement from the verified functional program presented here.

1.5.1 The compiler and stack machine

The program is given in Figure 1.1, in two halves. In the first half we reiterate the definitions of the `IntExp` type and its evaluation function `eval`, which is defined by structural recursion over `IntExp`.

In the second half of the figure we give a model of a stack machine which is used to evaluate the expressions. The machine operates over a stack of integers, hence the definition

```
type Stack = [Int]
```

The instructions for the machine are given by the type `Code`, which has two operations, namely to push an element (`PushLit`) onto the stack and to perform an evaluation of an operation (`DoBinary`) using the top elements of the stack as arguments.

An expression is converted into a `Program`, that is a list of `Code`, by `compile`. The `compile` function compiles a literal in the obvious way, and for an operator expression, the compiled code consists of the compiled code for the two expressions, concatenated by the list operator `++`, with the appropriate binary operator invocation appended.

The operation of the machine itself is described by

```
run :: Program -> Stack -> Stack
```

and from that definition it can be seen that if the stack fails to have at least two elements on operator evaluation, execution will be halted and the stack cleared.

1.5.2 Formulating the goal

The intended effect of the compiler is to produce code (for `e`) which when run puts the value of `e` on the stack. In formal terms,

$$\text{run (compile e) []} \equiv [\text{eval e}] \tag{10}$$

Now, we could look for a proof of this by structural induction over `e`, but this will fail. We can explain this failure from two different points of view.

Looking first at the problem itself, we can see that in fact the compiler and machine have a rather more general property: no matter what the initial configuration of the stack, the result of the run should be to place the value of the expression on the top of the stack:

$$\text{run (compile e) stack} \equiv (\text{eval e : stack}) \quad (11)$$

This is still not general enough, since it talks about complete computations – what if the code is followed by more program? The effect should be to evaluate e and place its result on the stack prior to executing the remaining program. We thus reach the final formulation of the goal

$$\begin{aligned} \text{run (compile e ++ program) stack} \\ \equiv \text{run program (eval e : stack)} \end{aligned} \quad (12)$$

An alternative view of the difficulty comes from looking at the failed proof attempt: the induction hypothesis turns out not to be powerful enough to give what is required. When this happens we can use the mismatch to find the appropriate generalisation of the hypothesis — the reader can try this for herself. A guide to the form of hypothesis is often given by the form taken by the definitions of the functions under scrutiny; we will illustrate this point in Section 1.5.3 below.

1.5.3 The proof

Our goal is to prove (12) for all values of e , program and stack . As a first attempt we might try to prove (12) by induction over e , for arbitrary program and stack , but this will fail. This happens because the induction hypothesis will be used at different values of stack and program , so that the goal for the inductive proof is to show by structural induction on e that

$$\begin{aligned} (\forall \text{program, stack}) (\text{run (compile e ++ program) stack} \\ \equiv \text{run program (eval e : stack)}) \end{aligned} \quad (13)$$

holds for all e .

The proof is given in Figure 1.2 and follows the principle of structural induction for IntExp presented in Section 1.4 above. In the first part we prove the base case:

$$\begin{aligned} \text{run (compile (Literal int) ++ program) stack} \\ \equiv \text{run program (eval (Literal int) : stack)} \end{aligned} \quad (14)$$

for arbitrary program, stack , thus giving the base case of (13). The proof proceeds by separately rewriting the left- and right-hand sides of (14) to the same value.

In the second part we show

$$\begin{aligned} \text{run (compile (Binary op ex1 ex2) ++ program) stack} \\ \equiv \text{run program (eval (Binary op ex1 ex2) : stack)} \end{aligned} \quad (15)$$

for arbitrary `program, stack` using the induction hypotheses for `ex1`:

$$\begin{aligned} & (\forall \text{program, stack}) (\text{run (compile ex1 ++ program) stack} \\ & \quad \equiv \text{run program (eval ex1 : stack)}) \end{aligned} \tag{16}$$

and `ex2`. It is instructive to observe that in the proof the induction hypothesis for `ex1`, (16), is used with the expression

```
compile ex2 ++ [DoBinary op] ++ program
```

substituted for `program`, and that for `ex2` is used in a similar way. Again the proof proceeds by separately rewriting the left- and right-hand sides of (15).

The third part of Figure 1.2 shows how our original goal, (10) is a consequence of the more general result (12).

How might we be led to the goal (13) by the form of the program itself? If we examine the definition of `run` we can see that in the recursive calls (7) and (8) the `stack` parameter is modified. This indicates that the `stack` cannot be expected to be a parameter of the proof, and so that the general formulation of the induction hypothesis will have to include all possible values of the stack parameter.

1.6 General recursion

In the preceding sections we saw how structural recursion and induction can be used to define and verify programs over algebraic data types. Functions defined in this way are manifestly total, but there remains the question of whether these limited forms of recursion and induction are adequate in practice. An example going beyond structural recursion over `IntExp` is a function to rearrange arithmetic expressions so that the additions which they contain are associated to the left, transforming

$$(4+2)+(3+(7+9)) \quad \text{to} \quad (((4+2)+3)+7)+9$$

The function is defined thus:

```
lAssoc :: IntExp -> IntExp

lAssoc (Literal n) = Literal n
lAssoc (Binary Sub ex1 ex2)
  = Binary Sub (lAssoc ex1) (lAssoc ex2)
lAssoc (Binary Add ex1 (Binary Add ex3 ex4))
  = lAssoc (Binary Add (Binary Add ex1 ex3) ex4) \tag{17}
lAssoc (Binary Add ex1 ex2)
  = Binary Add (lAssoc ex1) (lAssoc ex2)
```

(where the `Mul` case has been omitted). Each clause is structurally recursive, except for (17), in which the top-level expression `ex1+(ex3+ex4)` is transformed to `(ex1+ex3)+ex4`. Once this transformation has been effected, it is necessary

to re-examine the whole re-arranged expression, and not just the components of the original. The reader might like to experiment with the example expression to convince herself of the necessity of making a definition of this form, rather than a structural recursion.

Now, what is the lesson of examples like this for the design of functional programming languages and for verification of systems written in them? There are broadly two schools of thought.

The predominant view is to accept that a language should allow arbitrary recursion in the definitions of functions (and perhaps other objects). Mainstream languages such as Haskell, Miranda and Standard ML are all of this kind. With arbitrary recursion come a number of consequences.

- The semantics of the language becomes more complex, since it must now contain an account of the possible non-termination of programs.
- Moreover, the evaluation mechanism becomes significant. If all programs terminate, then the order in which programs are evaluated is not an issue; if non-termination is possible then strict and lazy evaluation strategies differ, and thus give lazy and strict languages different semantics.
- As far as the topic of this chapter is concerned, the complexity of the semantics is reflected in the logic needed to reason about the language, for both strict and lazy languages.

For these reasons there has been recent interest in terminating languages — Turner’s notion of ‘strong’ functional languages [14] — because such languages both have a simpler proof theory and have full freedom of choice for evaluation strategy, which is of course of relevance to the field of parallel functional programming.

In the remainder of this chapter we will explore the effect of these two alternatives for functional program verification, first looking at the mainstream, partial, languages.

1.7 Partial languages

This section gives an informal overview of the effect of admitting general recursion into a programming language, and emphasises the consequent split between lazy and strict languages. This serves as an introduction to the overview of the semantic basis of languages with partiality in the section to come.

1.7.1 Strict languages

In a strict language such as (the pure subset of) Standard ML arbitrary forms of recursive definitions are allowed for functions. A definition of the form

```
undefFun :: t -> t
undefFun x = undefFun x
```

(18)

(using Haskell-style syntax) has the effect of forcing there to be an undefined element at every type. What effect does this have for evaluation and for the logic? Take the example function

```
const :: s -> t -> t
const a b = a
```

and consider its logical translation. Our earlier work suggests that we translate it as

$$\text{const } a \ b \equiv a \tag{19}$$

but we need to be careful what is substituted for the variables **a** and **b**. If we take **a** to be **3** and **b** to be **undefFun 4** then it appears that

$$\text{const } 3 \ (\text{undefFun } 4) \equiv 3$$

This is contrary to the rule for evaluation which states that arguments need to be evaluated prior being passed to functions, and which means that (18) should be undefined when applied to **undefFun 4**. The translation (19) can therefore only apply to **values** (of type **Int**) rather than arbitrary **expressions** of that type as was the case earlier. This can be made clear by re-expressing (19) thus:

$$(\forall_v a, b) (\text{const } a \ b \equiv a) \tag{20}$$

where the subscript in the quantifier ' \forall_v ' serves as a reminder that the quantifier ranges over all (defined) values rather than all expressions including those which denote an undefined computation.

1.7.2 Lazy languages

In a lazy language like Haskell or Miranda the definition of **undefFun** in (18) also gives rise to an undefined element at each type. This does not however affect the translation of **const** given in (19) above, since in a lazy language expressions are passed *unevaluated* to functions. In other words, the evaluation mechanism can truly be seen to be one of substitution of expressions for expressions. (For efficiency, this 'call by name' strategy will be implemented by a 'call by need' discipline under which the results of computations are shared.)

Nevertheless, the presence of an undefined expression in each type has its effect. We accept as a law the assertion that for all integers **x**

$$x+1 > x$$

but this will not be the case if **x** is an undefined computation. We will therefore have to make the distinction between defined values and all expressions as in Section 1.7.1.

The result of combining lazy evaluation and general recursion are more profound than for a strict language, since data structures can become partial or infinite. The effect of

```

nums = from 1
from n = n : from (n+1)

```

is to define the infinite list of positive integers, $[1, 2, 3, \dots]$. If `nums` is passed to a function, then it is substituted unevaluated, and parts of it are evaluated when and if they are required:

```

sft :: [Int] -> Int
sft (a:b:_) = a+b

```

(21)

```

sft nums
=> sft (from 1)
=> sft (1 : from 2)
=> sft (1 : 2 : from 3)

```

At this point the pattern match in can be performed, giving the result 3. Our interpretation therefore needs to include such infinite lists, as well as ‘partial’ lists such as `(2:undefFun 2)`. Note that under a strict interpretation all infinite and partial lists are identified with the undefined list, since they all lead to non-terminating computations.

In order to give a proper account of the behaviour of languages with non-termination we now look at the ways in which a formal or mathematical semantics can be given to a programming language.

1.8 Semantic approaches

This section surveys the two semantic approaches to functional programming languages with the aim of motivating the logical rules to which the semantics lead.

1.8.1 Denotational semantics

Under a denotational semantics, the objects of a programming language — both terminating and non-terminating — are modelled by the elements of a **domain**. A domain is a partially ordered structure, where the partial order reflects the degree of definedness of the elements, with the totally undefined object, \perp , lying below everything: $\perp \sqsubseteq x$. Recursion, as in the definition

$$f = C[f]$$

can then be explained by first looking at the sequence of approximations, f_n , with

$$f_0 \equiv \perp$$

and

$$f_{n+1} = C[f_n]$$

A domain also carries a notion of limit for a sequences (or indeed more general ‘directed sets’), so that the meaning of \mathbf{f} , $\llbracket \mathbf{f} \rrbracket$, is taken to be the limit of this sequence of approximations:

$$\mathbf{f} \equiv \bigsqcup_{\mathbf{n}} \mathbf{f}_{\mathbf{n}}$$

Another way of seeing this is that $\llbracket \mathbf{f} \rrbracket$ is the **least fixed point** of the operation $\lambda \mathbf{f} . \mathbf{C}[\mathbf{f}]$

with a domain having sufficient structure to provide fixed points of (monotone) operators over them.

All the data types of a functional language can be modelled in such a way, and reasoning over domains is characterised by fixed-point induction, which captures the fact that a recursively defined function is the limit of a sequence.

Fixed-point induction. If P is an inclusive predicate and if \mathbf{f} is defined as above, then if

$$\bullet P(\perp) \text{ holds, and} \tag{22}$$

$$\bullet P(\mathbf{f}_{\mathbf{n}}) \text{ implies } P(\mathbf{f}_{\mathbf{n}+1}); \tag{23}$$

then P holds of the limit of the sequence, that is $P(\mathbf{f})$.

A predicate P is inclusive if it is closed under taking limits, broadly speaking. Winskel, [15], provides a more detailed characterisation of this, together with sufficient conditions for a formula to be an inclusive predicate.

As an example we look again at the `lAssoc` function, defined in Section 1.6 above. We would like to show that rearranging an expression will not change its value, that is

$$P_0(\text{lAssoc}): \quad (\forall e) (\text{eval } (\text{lAssoc } e) \equiv \text{eval } e)$$

(where `eval` is defined in Section 1.4). Equations are inclusive, but unfortunately we cannot prove the inductive goals (22) and (23) for this particular property. We can modify it to say that *if the result is defined* then the equality holds, namely,

$$P(\text{lAssoc}): \quad (\forall e) ((\text{lAssoc } e \equiv \perp) \vee \text{eval } (\text{lAssoc } e) \equiv \text{eval } e)$$

It is interesting to see that this is a **partial correctness** property, predicated on the termination of the `lAssoc` function, for which we have to prove a separate termination result. We discuss this presently. To establish this result we have to prove (22) and (23) for this property. A proof of (22) is straightforward, since $P(\perp)$ states:

$$(\forall e) ((\perp e \equiv \perp) \vee \text{eval } (\text{lAssoc } e) \equiv \text{eval } e)$$

and clearly $\perp e \equiv \perp$ holds. A proof of (23) requires that we show that $P(\text{lAssoc}_{\mathbf{n}})$ implies $P(\text{lAssoc}_{\mathbf{n}+1})$ where (omitting the `Mul` case),

$$\text{lAssoc}_{n+1} \text{ (Literal } n) = \text{Literal } n \quad (24)$$

$$\begin{aligned} \text{lAssoc}_{n+1} \text{ (Binary Sub } ex1 \text{ } ex2) \\ = \text{Binary Sub (lAssoc}_n \text{ } ex1) \text{ (lAssoc}_n \text{ } ex2) \end{aligned} \quad (25)$$

$$\begin{aligned} \text{lAssoc}_{n+1} \text{ (Binary Add } ex1 \text{ (Binary Add } ex3 \text{ } ex4)) \\ = \text{lAssoc}_n \text{ (Binary Add (Binary Add } ex1 \text{ } ex3) \text{ } ex4) \end{aligned} \quad (26)$$

$$\begin{aligned} \text{lAssoc}_{n+1} \text{ (Binary Add } ex1 \text{ } ex2) \\ = \text{Binary Add (lAssoc}_n \text{ } ex1) \text{ (lAssoc}_n \text{ } ex2) \end{aligned} \quad (27)$$

Now, our goal is to prove that

$$(\forall e)((\text{lAssoc}_{n+1} \text{ } e \equiv \perp) \vee \text{eval (lAssoc}_{n+1} \text{ } e) \equiv \text{eval } e)$$

on the assumption that

$$(\forall e)((\text{lAssoc}_n \text{ } e \equiv \perp) \vee \text{eval (lAssoc}_n \text{ } e) \equiv \text{eval } e)$$

We look at the cases of the definition in turn. For a literal we have by (24)

$$\text{lAssoc}_{n+1} \text{ (Literal } n) \equiv \text{Literal } n$$

from which we conclude immediately that

$$\text{eval (lAssoc}_{n+1} \text{ (Literal } n)) \equiv \text{eval (Literal } n)$$

Now, looking at subtraction, and assuming that the function terminates, we have

$$\begin{aligned} \text{eval (lAssoc}_{n+1} \text{ (Binary Sub } ex1 \text{ } ex2)) \\ \equiv \{ \text{by (25)} \} \\ \text{eval (Binary Sub (lAssoc}_n \text{ } ex1) \text{ (lAssoc}_n \text{ } ex2)) \\ \equiv \{ \text{by definition of } \text{eval} \} \\ \text{eval (lAssoc}_n \text{ } ex1) - \text{eval (lAssoc}_n \text{ } ex2) \\ \equiv \{ \text{by termination and the induction hypothesis} \} \\ \text{eval } ex1 - \text{eval } ex2 \\ \equiv \{ \text{by definition of } \text{eval} \} \\ \text{eval (Binary Sub } ex1 \text{ } ex2) \end{aligned}$$

The tricky case is (26), which is the non-structurally recursive clause. Now, again assuming termination, we have

$$\begin{aligned} \text{eval (lAssoc}_{n+1} \text{ (Binary Add } ex1 \text{ (Binary Add } ex3 \text{ } ex4))) \\ \equiv \{ \text{by (26)} \} \\ \text{eval (lAssoc}_n \text{ (Binary Add (Binary Add } ex1 \text{ } ex3) \text{ } ex4)) \\ \equiv \{ \text{by termination and the induction hypothesis} \} \\ \text{eval ((Binary Add (Binary Add } ex1 \text{ } ex3) \text{ } ex4)) \\ \equiv \{ \text{by the associativity of } + \} \\ \text{eval (Binary Add } ex1 \text{ (Binary Add } ex3 \text{ } ex4)) \end{aligned}$$

We leave the final case as an exercise. This establishes the induction step, and so the result itself.

How do we prove that `lAssoc` terminates on all arguments? We need to have some ‘measure of progress’ in the recursive calls. In all calls but (17) the recursive calls are on structurally smaller expressions, but in (17) the call is to an expression containing the same number of operators. What is changed in the recursive call is the arrangement of the expression, and it is easy to see that on the right hand side of the `Add` in the recursive call there are fewer applications of `Add` than in the same position on the left hand side:

$$\begin{array}{ccc}
 & + & \\
 & / \ \backslash & \\
 e1 & + & \\
 & / \ \backslash & \\
 & e2 \ e3 & \\
 & & \\
 & & + \\
 & & / \ \backslash \\
 & & + \ e3 \\
 & & / \ \backslash \\
 & & e1 \ e2
 \end{array}$$

This reduction means that there can only be a finite number of repeated calls to (17) before one of the structural cases is used. Informally, what we have done is to give an ordering over the expressions which is **well-founded**, that is has no infinite descending chains (like the chain $-1 > -2 > \dots > -n > \dots$) over the integers. A recursion will terminate precisely when it can be shown to follow a well-founded ordering.

Further details about denotational semantics can be found in [15, 6]. We also refer back to denotational semantics at the end of section 1.8.3

1.8.2 Operational semantics

The structured (‘SOS’) style of operational semantics pioneered by Plotkin describes a programming language by means of deduction rules which explain how expressions are evaluated. This style has been used to describe real languages, notably Standard ML [4], and arguably it gives a more readable and concise description of a language than a denotational semantics. The account given in this section relies on Gordon’s thesis, [2], which applies these ideas to the description of functional programming languages.

SOS descriptions give reduction rules (describing ‘one step’ of the computation), as in

$$((\lambda x.M)N) \longrightarrow M[N/x]$$

or can provide a description of the evaluation of an expression to a value (the ‘big step’ rules), thus:

$$\frac{L \Longrightarrow (\lambda x.M) \quad M[N/x] \Longrightarrow V}{(L N) \Longrightarrow V}$$

These rules are related, with \Longrightarrow representing arbitrarily many steps under the relation \longrightarrow . From these rules an equality relation can be generated: two expressions are equal, $L \simeq M$, if whatever context $C[_]$ they are placed in, $C[L] \Longrightarrow V$ if and only if $C[M] \Longrightarrow V$. Now, the issue becomes one of finding ways of deducing, for given expressions L and M , that $L \simeq M$ holds. Abramsky

had the insight that this relation resembled the bisimulations of process calculi. This characterises the equivalence as a *greatest* fixed point. Rather than look at the general theory of bisimulations, we will look here at how it applies to infinite lists.

The equality relation over infinite lists, ‘ \simeq ’, is the greatest fixed point of the definition

$$l \simeq m \iff_{df} \text{there exist } a, b, l', m' \text{ so that } l \longrightarrow (a:l'), \\ m \longrightarrow (b:m'), a \equiv b \text{ and } l' \simeq m'.$$

Now, the greatest fixed point of a relation can be characterised as the union of all the post-fixed points of the relation, which in this case are called **bisimulations**. The relation \mathcal{S} is a bisimulation if

$$l \mathcal{S} m \implies \text{there exist } a, b, l', m' \text{ so that } l \longrightarrow (a:l'), \\ m \longrightarrow (b:m'), a \equiv_{\mathcal{S}} b \text{ and } l' \equiv_{\mathcal{S}} m'.$$

where $\equiv_{\mathcal{S}}$ is the smallest congruence generated by the relation \mathcal{S} . It is now the case that

$$l \simeq m \iff \text{there exists a bisimulation } \mathcal{S} \text{ such that } l \mathcal{S} m.$$

In the next section we give an example of a proof using this definition of bisimulation.

1.8.3 An example of coinduction

In this section we give proof of the equality of two lists of the factorials of the natural numbers. The first is a mapping of the factorial function along the list of natural numbers, while the second, `facts 0`, gives a recursive definition of the list in question.

```
facMap :: [Integer]
facMap = map fac [0..]
```

```
fac :: Integer -> Integer
fac 0 = 1
```

(32)

```
fac (n+1) = (n+1) * fac n
```

(33)

```
facs :: Integer -> [Integer]
facs n = fac n : zipWith (*) [(n+1)..] (facs n)
```

(34)

To prove the equality of the two lists `facMap` and `facs 0` we first prove an auxiliary result, namely that

$$\text{zipWith } (*) \text{ [(n+1)..] (facs n) } \simeq \text{facts (n+1)}$$
(35)

for all natural numbers n . In order to do this we take the relation

$$\mathcal{S} \equiv \{ (\text{zipWith } (*) \text{ [(n+1)..] (facs n) }, \text{facts (n+1)}) \mid n \in \text{Nat} \}$$

and show that it is a bisimulation. Expanding first the left hand side of a typical element we have

```
zipWith (*) [(n+1)..] (facs n)
  => zipWith (*) (n+1:[(n+2)..]) (fac n : (tail (facs n)))
  => (n+1)*(fac n) : zipWith (*) [n+2..] (zipWith (*) [(n+1)..] (facs n))
  => fac (n+1) : zipWith (*) [n+2..] (zipWith (*) [(n+1)..] (facs n))
```

On the right hand side we have

```
facs (n+1)
  => fac (n+1) : zipWith (*) [n+2..] (facs (n+1))
```

Now observe the two expressions. They have equal heads, and their tails are related by $\equiv_{\mathcal{S}}$ since they are applications of the function

```
zipWith (*) [(n+2)..]
```

to lists which are related by \mathcal{S} , namely

```
zipWith (*) [(n+1)..] (facs n)  $\simeq$  facs (n+1)
```

This establishes the result (35), and the consequence that

$$\text{facs } n \simeq \text{fac } n : \text{facs } (n+1) \tag{36}$$

Now we prove that

```
facs n  $\simeq$  map fac [n..]
```

by showing that the relation

$$\mathcal{R} \equiv \{ (\text{facs } n, \text{map fac } [n..]) \mid n \in \text{Nat} \}$$

is a bisimulation. Taking a typical pair, we have,

```
facs n                                map fac [n..]
   $\simeq$  fac n : facs (n+1)                => fac n : map f [(n+1)..]
```

which establishes that \mathcal{R} is a bisimulation and in particular shows that

```
facs 0  $\simeq$  map fac [0..]
```

as we sought.

It is interesting to observe that recent work has shown that coinduction principles can be derived directly in domain theory; see [9] for more details.

1.9 Strong functional programming

We have seen that the potential for non-termination makes program verification more complicated. Because of this and other reasons there is interest in programming languages which are ‘strong’ in the sense of providing only the means to define terminating functions. In this section we give a brief overview of various of these research directions. A general point to examine is the degree to which each approach limits a programmer’s expressivity.

1.9.1 Elementary strong functional programming

Turner [14] proposes a language with limited recursion and co-recursion as a terminating functional language which could be used by beginning programmers (in contrast to alternatives discussed later in this section). The language proposed will have compile-time checks for the termination of recursive definitions, along the lines of [3, 11]. The language also contains co-recursion, the dual of recursion, over co-data, such as infinite lists (the greatest fixed point of a particular type equality). The definition

```
facts = 1 : zipWith (*) [1..] facts
```

is recognisable as a ‘productive’ definition, since the recursive call to `facts` on the right hand side is protected within the constructor ‘:’. Note the duality with primitive recursion, like

```
length (a:x) = 1 + length x
```

in which the recursive call to `length` is on a component of the argument, `(a:x)`, which is contained in the application of the constructor ‘:’. Proof of properties of these corecursive objects is by coinduction, as discussed above.

The disadvantage of this approach is that it must rely on the compile-time algorithms which check for termination. It is not clear, for instance, whether the earlier definition of the `lAssoc` function is permitted in this system, and so the expressivity of the programmer is indeed limited by this approach. On the other hand, it would be possible to implement such a system as a ‘strong’ subset of an existing language such as Haskell, and to gain the advantage of remaining in the terminating part of the language whenever possible.

1.9.2 Constructive type theories

Turner’s language eschews the more complex dependent types of the constructive type theories of Martin-Löf and others [5, 12]. These languages are simultaneously terminating functional languages and constructive predicate logics, under the Curry/Howard Isomorphism which makes the following identifications:

Programming		Logic
Type		Formula
Program		Proof
Product/record type	&	Conjunction
Sum/union type	∨	Disjunction
Function type	->	Implication
Dependent function type	∀	Universal quantifier
Dependent product type	∃	Existential quantifier
...		...

in which it is possible in an integrated manner to develop programs and their proofs of correctness.

From the programming point of view, there is the addition of dependent types, which can be given by functions which return different types for different argument values: an example is the type of vectors, `Vec`, where `Vec(n)` is the type of vectors of length `n`. Predicates are constructed in a similar way, since a predicate yields different logical propositions – that is types – for different values.

Predicates (that is dependent types) can be constructed inductively as a generalisation of algebraic types. We might define the less than predicate ‘<’ over `Nat` by saying that there are two constructors for the type:

```
ZeroLess :: (∀n::Nat) (0 < S n)
SuccLess :: (∀n::Nat) (∀m::Nat) ((m < n) -> (S m < S n))
```

This approach leads to a powerful style of proof in which inductions are performed over the form of proof objects, that is the elements of types like `(m < n)`, rather than over (say) the natural numbers, and such a method makes much more manageable a proof of the transitivity of ‘<’ over `Nat`, say.

A more expressive type system allows programmers to give more accurate types to common functions, such as function which indexes the elements of a list.

```
index :: (∀l::[t]) (∀n::Nat) ((n < length l) -> t)
```

An application of `index` has *three* arguments: a list, `l` and a natural number `n` — as for the standard `index` function — and a third argument which is of type `(n < length l)`, that is a *proof* that `n` is a legitimate index for the list in question. This extra argument becomes a **proof obligation** which must be discharged when the function is applied to elements `l` and `n`.

The expressivity of a constructive type theory is determined by its proof-theoretic strength, so that a simple type theoretic language (without universes) would allow the definition of all functions which can be proved to be total in Peano Arithmetic, for instance. This includes most functions, except an interpreter for the language itself.

For further discussions of constructive type theories see [5, 12].

1.9.3 Program Transformation, Categories and Allegories

The histories of functional programming and program transformation have been intertwined from their inception. Serious program manipulations are not feasible in modern imperative languages which allow aliasing, pointer and reference modifications, type casting and so forth. More suited are current functional languages which support the definition of general operations – such as `map`, `filter` and `fold` over lists – as polymorphic higher-order functions. The properties of these functions – such as `map (f.g) = map f . map g` – can be expressed in a logic which extends the definitional equality of the programming language, and largely equational reasoning in that logic allows transformations to be written down in a formal way.

The most developed example of this work is Bird and de Moor's [1] in which they use the constructs of category theory to express their functional programming language.

Their categorical approach means that they are able to provide general rules for equational program manipulation. Prominent among these is the Fusion Law which states that

$$h \cdot (| f |) \equiv (| g |)$$

in the circumstances that

$$h \cdot f \equiv g \cdot F h$$

The 'banana' brackets ($|\dots|$) denote a catamorphism, that is a generalisation of `foldr` over lists, and so the rule gives a situation in which a composition including a fold can be made into a fold itself. Moreover, the law applies uniformly to all algebraic initial data types, so that there is no need separately to develop theories for lists, binary trees, rose trees and so on.

A pleasant feature of the categorical approach is the degree to which their reasoning can be equational. In particular the McCarthy conditional form (which is the function-level equivalent of `if ...then ...else ...`) gives case-free reasoning for functions for which a more traditional approach would require proof by cases. It should however be observed that even simple programs like factorial require some manipulation to be put into a catamorphic form, and a two argument function like concatenation of two lists requires substantial work to put it into this form form.

In the second half of [1] Bird and de Moor replace the calculus of functions with a calculus of relations. This allows more freedom in specification, allowing systems to be specified as the inverse of a simple operation, or as the meet of two requirements, for instance. Another advantage of relations over functions is that the relations include the non-deterministic functions. This means that an algebra of relational programming allows us to reason about non-deterministic functions.

Finally, the relations also include the *partial* functions. So an algebra of relational programming provides all the familiar partial functions, but without the complexity of partial elements. In particular, it provides the fixpoint operator without the complexity of partial elements and divergence normally associated with it: a non-well-founded recursion will not produce a non-terminating function, but rather an empty relation. Note that this means that one disadvantage of the functional approach, namely the restriction to recursive functions that are catamorphisms, no longer applies in the relation setting.

A review of [1] which expands upon this section can be found in [10].

Bibliography

- [1] Richard Bird and Oege de Moor. *Algebra of Programming*. Prentice-Hall, 1997.
- [2] Andrew J. Gordon. *Functional Programming and Input/Output*. British Computer Society Distinguished Dissertations in Computer Science. Cambridge University Press, 1994.
- [3] D. McAllester and K. Arkondas. Walther recursion. In M.A. Robbie and J.K. Slaney, editors, *CADE 13*. Springer-Verlag, 1996.
- [4] Robin Milner, Mads Tofte, and Robert Harper. *The Definition of Standard ML*. MIT Press, 1990.
- [5] Bengt Nordström, Kent Petersson, and Jan M. Smith. *Programming in Martin-Löf's Type Theory — An Introduction*, volume 7 of *International Series of Monographs on Computer Science*. Oxford University Press, 1990.
- [6] Laurence C. Paulson. *Logic and Computation — Interactive proof with Cambridge LCF*. Cambridge University Press, 1987.
- [7] John Peterson and Kevin Hammond, editors. *Report on the Programming Language Haskell, Version 1.4*. www.haskell.org/report, 1997.
- [8] A. M. Pitts and I. D. B. Stark. Operational reasoning for functions with local state. In A. D. Gordon and A. M. Pitts, editors, *Higher Order Operational Techniques in Semantics*, Publications of the Newton Institute. Cambridge University Press, 1998.
- [9] Andrew M. Pitts. A co-induction principle for recursively defined domains. *Journal of Logic and Computation*, 1992.
- [10] Erik Poll and Simon Thompson. Review of [1]. *Journal of Functional Programming*, to appear.
- [11] Alastair Telford and David Turner. Ensuring streams flow. In M. Johnson, editor, *Algebraic Methodology and Software Technology 1997*. Springer-Verlag, 1997.

- [12] Simon Thompson. *Type Theory and Functional Programming*. Addison Wesley, 1991.
- [13] Simon Thompson. A Logic for Miranda, Revisited. *Formal Aspects of Computing*, 7, 1995.
- [14] David Turner. Elementary strong functional programming. In Pieter Hartel and Rinus Plasmeijer, editors, *Functional programming languages in education (FPLE)*, LNCS 1022. Springer-Verlag, Heidelberg, 1995.
- [15] Glynn Winskel. *The Formal Semantics of Programming Languages*. MIT Press, 1993.