

# The Conception, Evolution, and Application of Functional Programming Languages

Paul Hudak  
Yale University  
Department of Computer Science

March 1989  
Final Draft\*

## Abstract

The foundations of functional programming languages are examined from both historical and technical perspectives. Their evolution is traced through several critical periods: early work on lambda calculus and combinatory calculus, Lisp, Iswim, FP, ML, and modern functional languages such as Miranda<sup>1</sup> and Haskell. The fundamental premises on which the functional programming methodology stands are critically analyzed with respect to philosophical, theoretical, and pragmatic concerns. Particular attention is paid to the main features that characterize modern functional languages: higher-order functions, lazy evaluation, equations and pattern-matching, strong static typing and type inference, and data abstraction. In addition, current research areas—such as parallelism, non-determinism, input/output, and state-oriented computations—are examined with the goal of predicting the future development and application of functional languages.

**Categories and Subject Descriptors:** D.1.1 [Programming Techniques]: Applicative (Functional) Programming; D.3.2 [Programming Languages]: Language Classifications—applicative languages; data-flow languages; nonprocedural languages; very high-level languages; F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—lambda calculus and related systems; K.2 [History of Computing]: Software.

**General Terms:** Languages.

**Additional Key Words and Phrases:** referential transparency, higher-order functions, types, lazy evaluation, data abstraction.

## 1 Introduction

The earliest programming languages were developed with one simple goal in mind: to provide a vehicle through which one could control the behavior of computers. Not

---

\*To appear in ACM Computing Surveys.

<sup>1</sup>Miranda is a trademark of Research Software Ltd.

surprisingly, these early languages reflected fairly well the structure of the underlying machines. Although at first blush this seems eminently reasonable, this viewpoint quickly changed, for two very good reasons. First, it became obvious that what was easy for a machine to reason about was not necessarily easy for a human to reason about. Second, as the number of different kinds of machines increased, the need arose for a common language with which to program all of them.

Thus from primitive assembly languages (which were at least a step up from raw machine code) there grew a plethora of high-level programming languages, beginning with Fortran in the 50's. The development of these languages grew so rapidly that by the 80's they were best characterized by grouping them into "families" that reflected a common computation model or programming "style." Debates over which language or family of languages is best will undoubtedly persist for as long as computers need programmers.

The class of *functional*, or *applicative* programming languages, in which computation is carried out entirely through the *evaluation of expressions*, is one such family of languages, and debates over its merits have been quite lively in recent years. Are functional languages toys, or are they tools? Are they artifacts of theoretical fantasy, or of visionary pragmatism? Will they ameliorate our software woes, or merely compound them? Whatever answers one might have for these questions, one cannot ignore the significant interest current researchers have in functional languages, and the impact they have had on both the theory and pragmatics of programming languages in general.

Among the claims made by functional language advocates are that programs can be written quicker, are more concise, are "higher level" (resembling more closely traditional mathematical notation), are more amenable to formal reasoning and analysis, and can be executed more easily on parallel architectures. Of course, many of these features touch on rather subjective issues, which is one reason why the debates can be so lively.

It is my purpose in writing this paper to give the reader significant insight into the very essence of functional languages and the programming methodology that they support. I will do so by starting with a general discussion of the nature of functional languages, followed by a historical sketch of their development, a summary of the distinguishing characteristics of modern functional languages, and a discussion of current research areas. Through this study I hope to put into perspective both the power and weaknesses of the functional programming paradigm.

**A Note to the Reader:** I have assumed in writing this paper that the reader has a good understanding of the fundamental issues in programming language design and use. Although reasonably self-contained, this paper should not be considered as a tutorial on functional languages; the reader is referred to either [BW88] or [FH88] as suitable textbooks for learning more about modern functional programming techniques, including the important ideas behind *reasoning* about functional programs. In addition, I will say very little about how to *implement* functional languages; see [PJ87] for

an introduction to this important area (additional references are given in Sections 2.8 and 6).

Finally, a comment on notation: unless otherwise stated, all examples will be written in Haskell, a recently proposed functional language standard [HWe88]. However, rather than explain the syntax all at once, explanations will be given “on the fly” and enclosed in square brackets [...].<sup>2</sup>

## 1.1 The Programming Language Spectrum

*Imperative languages* are characterized as having an *implicit state* that is modified (i.e. side-effected) by constructs (i.e. commands) in the source language. As a result, such languages generally have a notion of *sequencing* (of the commands) to permit precise and deterministic control over the state. Most, including the most popular, languages in existence today are imperative.

As an example, the *assignment statement* is a (very common) *command*, since its effect is to alter the underlying implicit store so as to yield a different binding for a particular variable. The `begin . . . end` construct is the prototypical *sequencer* of commands, as are the well-known *goto statement* (unconditional transfer of control), *conditional statement* (qualified sequencer), and *while loop* (an example of a “structured command”). With these simple forms we can then, for example, compute the factorial of the number `x`:

```
n := x;
a := 1;
while n>0 do
begin a := a*n;
      n := n-1
end;
...
```

After execution of this program the value of `a` in the implicit store will contain the desired result.

In contrast, *declarative languages* are characterized as having *no* implicit state, and thus the emphasis is placed entirely on programming with *expressions* (or *terms*). In particular, *functional languages* are declarative languages whose underlying model of computation is the *function* (in contrast to, for example, the *relation* that forms the basis for logic programming languages).

---

<sup>2</sup>Since the Haskell Report is relatively new, some minor changes to the language may occur after this paper has appeared. An up-to-date copy of the Report may be obtained from the author.

In a declarative language state-oriented computations are accomplished by carrying the state around explicitly rather than implicitly, and looping is accomplished via *recursion* rather than by sequencing. For example, the factorial of  $x$  may be computed in the functional language Haskell by:

```
fac x = 1
where fac n a = if n>0 then fac (n-1) (a*n)
              else a
```

in which the formal parameters  $n$  and  $a$  are examples of “carrying the state around explicitly,” and the recursive structure has been arranged so as to mimic as closely as possible the looping behavior of the program given earlier. Note that the conditional in this program is an *expression* rather than *command*; i.e. it denotes a *value* (conditional on the value of the predicate), rather than denoting a sequencer of commands. Indeed the *value of the program* is the desired factorial, rather than it being found in an implicit store.

Functional (in general, declarative) programming is often described as expressing *what* is being computed rather than *how*, although this is really a matter of degree. For example, the above program may say less about how factorial is computed than the imperative program given earlier, but is perhaps not as abstract as:

```
fac x = 1
where fac n = if n==0 then 1
              else n * fac (n-1)
```

[`==` is the infix operator for equality.] which appears very much like the mathematical definition of factorial, and is indeed a valid functional program.

Since most languages have expressions, it is tempting to take our definitions literally and describe functional languages via derivation from conventional programming languages: simply drop the assignment statement and any other side-effecting primitives, and there you have it! This, of course, is very misleading. The result of such a derivation is usually far less than satisfactory, since the purely functional subset of most imperative languages is hopelessly weak (although there are important exceptions, such as Scheme [RCe86]).

Rather than saying, then, what functional languages *don't* have, it is better to characterize them by the features they *do* have. For modern functional languages, those features include higher-order functions, lazy evaluation, pattern-matching, and various kinds of data abstraction — all of these features will be described in detail in this paper. Functions are treated as first-class objects, are allowed to be recursive, higher-order, and polymorphic, and in general are provided with mechanisms that ease their definition and use. Syntactically, modern functional languages have an “equational” look, in which functions are defined using mutually recursive equations and pattern-matching.

This discussion suggests that what is important is the functional programming *style*, in which the above features are manifest, and in which side effects are strongly discouraged but not necessarily eliminated. This is the viewpoint taken by the ML community, for example, and to some extent the Scheme community. On the other hand, there is a very large contingency of “purists” in the functional programming community who believe that purely functional languages are not only sufficient for general computing needs, but also better because of their “purity.” At least a dozen purely functional languages exist along with their implementations.<sup>3</sup> The main property that is lost when side effects are introduced is *referential transparency*; this loss in turn impairs *equational reasoning*, as described below.

## 1.2 Referential Transparency and Equational Reasoning

The emphasis on a pure declarative style of programming is perhaps the hallmark of the functional programming paradigm. The term *referentially transparent* is often used to describe this style of programming, in which “equals can be replaced by equals.” For example, consider the (Haskell) expression:

```
... x+x ...  
where x = f a
```

The function application (f a) may be substituted for any free occurrence of x in the scope created by the where expression, such as in the subexpression x+x. The same cannot generally be said of an imperative language, where one must first be sure that no assignment to x is made in any of the statements intervening between the initial definition of x and one of its subsequent uses.<sup>4</sup> In general this can be quite a tricky task, for example in the case where procedures are allowed to induce non-local changes to lexically scoped variables.

Although the notion of referential transparency may seem like a simple idea, the clean equational reasoning that it allows is very powerful, not only for reasoning formally about programs, but also informally in writing and debugging programs. A program in which side-effects are minimized, but not eliminated, may still benefit from equational reasoning, although naturally more care must be taken when applying such reasoning. The degree of care, however, may be much higher than one might think at first: most languages that allow “minor” forms of side-effects do not minimize their

---

<sup>3</sup>This situation forms an interesting contrast with the logic programming community, where Prolog is often described as declarative (whereas Lisp is usually not), and there are very few pure logic programming languages (and even fewer implementations).

<sup>4</sup>In all fairness, there are logics for reasoning about imperative programs, such as those espoused by Floyd, Hoare, Dijkstra, Wirth, and others. None of them, however, exploit any notion of referential transparency.

locality lexically—thus any call to any function in any module might conceivably introduce a side-effect, in turn invalidating many applications of equational reasoning.

The perils of side-effects are appreciated by the most experienced programmers in any language, although most are loathe to give them up completely. It remains the goal of the functional programming community to demonstrate that one can do *completely* without side-effects, without sacrificing efficiency or modularity. Of course, as mentioned earlier, the lack of side-effects is not all there is to the functional programming paradigm—as we shall soon see, modern functional languages rely heavily on certain other features, most notably higher-order functions, lazy evaluation, and data abstraction.

### 1.3 Plan of Study

Unlike many developments in computer science, functional languages have maintained the principles on which they were founded to a surprising degree. Rather than changing or compromising those ideas, modern functional languages are best classified as embellishments of a certain set of ideals. It is a distinguishing feature of modern functional languages that they have so effectively held on to pure mathematical principles in a way shared by very few other languages.

Because of this, one can learn a great deal about functional languages simply by studying their evolution closely. On the other hand, such a study may fail to yield a consistent treatment of any one feature that is common to most functional languages, for it will be fractured into its manifestations in each of the languages as they were historically developed. For this reason I have taken a three-fold approach to our study:

First, in the next section I will provide a historical sketch of the development of functional languages. Starting with the lambda calculus as the prototypical functional language, I will gradually embellish it with ideas as they were historically developed, leading eventually to a reasonable technical characterization of modern functional languages.

In so doing, however, I will postpone detailed discussion of four important concepts—higher-order functions, lazy evaluation, data abstraction mechanisms, and equations/pattern-matching—which are critical components of all modern functional languages and are best discussed as independent topics. Section 3 is devoted to these important ideas.

Finally, in Section 4 I will discuss more advanced ideas, and outline some critical research areas. Then to round out the paper, in Section 5 I will try to put into perspective some of the limitations of functional languages by examining some of the myths that have accompanied their development.

## 2 The Evolution of Functional Languages

### 2.1 The Lambda Calculus

The development of functional languages has been influenced from time to time by many sources, but none is as paramount nor as fundamental as the work of Alonzo Church on the *lambda calculus* [Chu33, Chu41]. Indeed the lambda calculus is usually regarded as the first functional language, although it was certainly not thought of as programming language at the time, given that there were no computers to run the programs on! In any case, modern functional languages can be thought of as (non-trivial) embellishments of the lambda calculus.

It is often thought that the lambda calculus also formed the foundation for Lisp, but this in fact appears not to be the case [McC78]. The impact of the lambda calculus on early Lisp development was minimal, and it has only been very recently that Lisp has begun to evolve more toward lambda calculus ideals. On the other hand, Lisp had a significant impact on the subsequent development of functional languages, as will be discussed in Section 2.2.

Church's work was motivated by the desire to create a *calculus* (informally, a syntax for terms and set of rewrite rules for transforming terms) that captured one's intuition about the behavior of *functions*. This approach is counter to the consideration of functions as, for example, *sets* (more precisely, sets of argument/value pairs), since the intent was to capture the *computational* aspects of functions. A calculus is a formal way for doing just that.

Church's *lambda calculus* was the first suitable treatment of the computational aspects of functions. Its type-free nature yielded a particularly small and simple calculus, and it had one very interesting property, capturing functions in their fullest generality: functions could be applied to themselves. In most reasonable theories of functions as sets, this is impossible, since it requires the notion of a set containing itself, resulting in well-known paradoxes. This ability of self-application is what gives the lambda calculus its power, since it allows one to gain the effect of recursion without explicitly writing a recursive definition. Despite this powerful ability, the lambda calculus is *consistent* as a mathematical system—no contradictions or paradoxes arise.

Because of the relative importance of the lambda calculus to the development of functional languages, I will describe it in some detail in the remainder of this section, using modern notational conventions.

#### 2.1.1 Pure Untyped Lambda Calculus

The abstract syntax of the *pure untyped lambda calculus* (a name chosen to distinguish it from other versions developed later) embodies what are called *lambda expressions*,

defined by:<sup>5</sup>

$$\begin{array}{ll}
 x \in Id & \text{Identifiers} \\
 e \in Exp & \text{Lambda Expressions} \\
 \text{where } e ::= x \mid e_1 e_2 \mid \lambda x.e &
 \end{array}$$

Expressions of the form  $\lambda x.e$  are called *abstractions*, and of the form  $(e_1 e_2)$  are called *applications*. It is the former that captures the notion of a function, and the latter that captures the notion of application of a function. By convention, application is assumed to be left associative, so that  $(e_1 e_2 e_3)$  is the same as  $((e_1 e_2) e_3)$ .

The rewrite rules of the lambda calculus depend on the notion of *substitution* of an expression  $e_1$  for all “free” occurrences of an identifier  $x$  in an expression  $e_2$ , which we write as  $[e_1/x]e_2$ .<sup>6</sup> Most systems, including both the lambda calculus and predicate calculus, that utilize substitution on identifiers must be careful to avoid name conflicts. Thus, although the intuition behind substitution is strong, its formal definition can be somewhat tedious—the reader who is comfortable with his/her intuition about substitution may skip over the next paragraph, which is included primarily for completeness.

To understand substitution we must first understand the notion of the *free variables* of an expression  $e$ , which we write as  $fv(e)$ , and define by the following simple rules:

$$\begin{array}{ll}
 fv(x) & = \{x\} \\
 fv(e_1 e_2) & = fv(e_1) \cup fv(e_2) \\
 fv(\lambda x.e) & = fv(e) - \{x\}
 \end{array}$$

We say that  $x$  is *free in*  $e$  iff  $x \in fv(e)$ . The substitution  $[e_1/x]e_2$  is then defined inductively by:

$$\begin{array}{ll}
 [e/x_i]x_j & = \begin{cases} e, & \text{if } i = j \\ x_j, & \text{if } i \neq j \end{cases} \\
 [e_1/x](e_2 e_3) & = ([e_1/x]e_2) ([e_1/x]e_3) \\
 [e_1/x_i](\lambda x_j.e_2) & = \begin{cases} \lambda x_j.e_2, & \text{if } i = j \\ \lambda x_j.[e_1/x_i]e_2, & \text{if } i \neq j \text{ and } x_j \notin fv(e_1) \\ \lambda x_k.[e_1/x_i]([x_k/x_j]e_2), & \text{otherwise, where } k \neq i, k \neq j, \\ & \text{and } x_k \notin fv(e_1) \cup fv(e_2) \end{cases}
 \end{array}$$

The last rule is the subtle one, since it is where a name conflict could occur, and is resolved by making a name change. The following example demonstrates application

<sup>5</sup>The notation “ $d \in D$ ” means that  $d$  is a “typical element” of the set  $D$ , whose elements may be distinguished by subscripting. In the case of identifiers, we assume that each  $x_i$  is unique; i.e.,  $x_i \neq x_j$  if  $i \neq j$ . The notation “ $d ::= alt1 \mid alt2 \mid \dots \mid altn$ ” is standard BNF syntax.

<sup>6</sup>In denotational semantics the notation  $e[v/x]$  is used to denote the function  $e'$  that is just like  $e$  except that  $e' x = v$ . Our notation of placing the brackets in front of the expression is to emphasize that  $[v/x]e$  is a *syntactic* transformation on the expression  $e$  itself.



of all three rules:

$$[y/x]((\lambda y.x)(\lambda x.x)x) \equiv (\lambda z.y)(\lambda x.x)y$$

To complete the lambda calculus we define three simple rewrite rules on lambda expressions:

1.  $\alpha$ -conversion (“renaming”):  $\lambda x_i.e \Leftrightarrow \lambda x_j.[x_j/x_i]e$ , where  $x_j \notin \text{fv}(e)$ .
2.  $\beta$ -conversion (“application”):  $(\lambda x.e_1)e_2 \Leftrightarrow [e_2/x]e_1$ .
3.  $\eta$ -conversion:  $\lambda x.(e x) \Leftrightarrow e$ , if  $x \notin \text{fv}(e)$ .

These rules, together with the standard equivalence relation rules for reflexivity, symmetry, and transitivity, induce a theory of *convertibility* on the lambda calculus, which can be shown to be consistent as a mathematical system.<sup>7</sup> The well-known Church-Rosser Theorem [CR36] (actually two theorems) is what embodies the strongest form of consistency, and has to do with a notion of *reduction*, which is the same as convertibility but restricted so that  $\beta$ -conversion and  $\eta$ -conversion only happen in one direction:

1.  $\beta$ -reduction:  $(\lambda x.e_1)e_2 \Rightarrow [e_2/x]e_1$ .
2.  $\eta$ -reduction:  $\lambda x.(e x) \Rightarrow e$ , if  $x \notin \text{fv}(e)$ .

We write “ $e_1 \overset{*}{\Rightarrow} e_2$ ” if  $e_2$  can be derived from zero or more  $\beta$ - or  $\eta$ -reductions or  $\alpha$ -conversions; in other words  $\overset{*}{\Rightarrow}$  is the reflexive, transitive closure of  $\Rightarrow$  including  $\alpha$ -conversions. Similarly,  $\overset{*}{\Leftrightarrow}$  is the reflexive, transitive closure of  $\Leftrightarrow$ . In summary,  $\overset{*}{\Rightarrow}$  captures the notion of “reducibility,” and  $\overset{*}{\Leftrightarrow}$  captures the notion of “intra-convertibility.”

**Definition:** A lambda expression is in *normal form* if it cannot be further reduced using  $\beta$ - or  $\eta$ -reduction.

Note that some lambda expressions have *no* normal form, such as:

$$(\lambda x. (x x)) (\lambda x. (x x))$$

where the only possible reduction leads to an identical term, and thus the reduction process is non-terminating.

Nevertheless, the normal form appears to be an attractive canonical form for a term, has a clear sense of “finality” in a computational sense, and is what we intuitively think of as the *value* of an expression. Obviously we would like for that value to be unique, and we would like to be able to find it whenever it exists. The Church-Rosser Theorems give us positive results for both of these desires.

<sup>7</sup>The lambda calculus as we have defined it here is what Barendregt [Bar84] calls the  $\lambda K\eta$ -calculus, and is slightly more general than Church’s original  $\lambda K$ -calculus (which did not include  $\eta$ -conversion). Furthermore, Church originally showed the consistency of the  $\lambda I$ -calculus [Chu41], an even smaller subset (it only allowed abstraction of  $x$  from  $e$  if  $x$  was free in  $e$ ). We will ignore the subtle differences between these calculi here—our version is the one most often discussed in the literature on functional languages.

### 2.1.2 The Church-Rosser Theorems

**Church-Rosser Theorem I:** *If  $e_0 \approx e_1$  then there exists an  $e_2$  such that  $e_0 \rightarrow^* e_2$  and  $e_1 \rightarrow^* e_2$ .*<sup>8</sup>

In other words, if  $e_0$  and  $e_1$  are intra-convertible, then there exists a third term (possibly the same as  $e_0$  or  $e_1$ ) to which they can both be reduced.

**Corollary:** *No lambda expression can be converted to two distinct normal forms (ignoring difference due to  $\alpha$ -conversion).*

One consequence of this result is that how we arrive at the normal form does not matter; i.e. the order of evaluation is irrelevant (this has important consequences for parallel evaluation strategies). The question then arises as to whether or not it is always possible to find the normal form (assuming it exists). We begin with some definitions.

**Definition:** A *normal-order reduction* is a sequential reduction in which, whenever there is more than one reducible expression (called a *redex*), the left-most one is chosen first. In contrast, an *applicative-order reduction* is a sequential reduction in which the left-most *innermost* redex is chosen first.

**Church-Rosser Theorem II:** *If  $e_0 \rightarrow^* e_1$  and  $e_1$  is in normal form, then there exists a normal-order reduction from  $e_0$  to  $e_1$ .*

This is a very satisfying result; it says that if a normal form exists, we can always find it; i.e. just use normal-order reduction. To see why applicative-order reduction is not always adequate, consider the following example:

applicative-order: $(\lambda x. y) ((\lambda x. x x) (\lambda x. x x))$ $\Rightarrow (\lambda x. y) ((\lambda x. x x) (\lambda x. x x))$ $\Rightarrow \vdots$	normal-order: $(\lambda x. y) ((\lambda x. x x) (\lambda x. x x))$ $\Rightarrow y$
--	--

We will return to the trade-offs between normal- and applicative-order reduction in Section 3.2. For now we simply note that the strongest completeness and consistency results have been achieved with normal-order reduction.

In actuality, one of Church's (and others') motivations for developing the lambda calculus in the first place was to form a foundation for all of mathematics (in the way that, for example, set theory is claimed to provide such a foundation). Unfortunately, all attempts to extend the lambda calculus sufficiently to form such a foundation failed to yield a consistent theory. Church's original extended system was shown inconsistent by the Kleene-Rosser Paradox [KR35]; a simpler inconsistency proof is embodied in what is known as the Curry Paradox [Ros82]. The only consistent systems that have

---

<sup>8</sup>Church and Rosser's original proofs of their theorems are rather long, and many have tried to improve on them since. The shortest proof I am aware of for the first theorem is fairly recent, and aptly due to Rosser [Ros82].

been derived from the lambda calculus are much too weak to claim as a foundation for mathematics, and the problem remains open today.

These inconsistencies, although disappointing in a foundational sense, did not slow down research on the lambda calculus, which turned out to be quite a nice model of functions and of computation in general. The Church-Rosser Theorem was an extremely powerful consistency result for a computation model, and in fact rewrite systems completely different from the lambda calculus are often described as “possessing the Church-Rosser property” or even anthropomorphically as being “Church-Rosser.”

### 2.1.3 Recursion, $\lambda$ -Definability, and Church’s Thesis

Another nice property of the lambda calculus is embodied in the following theorem:

**Fixpoint Theorem:** *Every lambda expression  $e$  has a fixpoint  $e'$  such that  $(e e') \rightsquigarrow e'$ .*

**Proof:** Take  $e'$  to be  $(Ye)$  where  $Y$ , known as the *Y combinator*, is defined by:

$$Y \equiv \lambda f. (\lambda x. f(x x)) (\lambda x. f(x x))$$

Then we have:

$$\begin{aligned} (Ye) &= (\lambda x. e(x x)) (\lambda x. e(x x)) \\ &= e((\lambda x. e(x x)) (\lambda x. e(x x))) \\ &= e(Ye) \end{aligned}$$

This surprising theorem (and equally surprising simple proof) is what has earned  $Y$  the name “paradoxical combinator.” The theorem is quite significant—it means that any recursive function may be written non-recursively (and non-imperatively!). To see how, consider a recursive function  $f$  defined by:

$$f \equiv \dots f \dots$$

This could be rewritten as:

$$f \equiv (\lambda f. \dots f \dots) f$$

where note that the inner occurrence of  $f$  is now bound. This equation essentially says that  $f$  is a fixpoint of the lambda expression  $(\lambda f. \dots f \dots)$ —but that is exactly what  $Y$  computes for us, so we arrive at the following *non-recursive* definition for  $f$ :

$$f \equiv Y (\lambda f. \dots f \dots)$$

As a concrete example, the factorial function:

$$fac \equiv \lambda n. \text{if } (n = 0) \text{ then } 1 \text{ else } (n * fac(n - 1))$$

can be written non-recursively as:

$$fac \equiv Y(\lambda fac. \lambda n. \text{if } (n = 0) \text{ then } 1 \text{ else } (n * fac(n - 1)))$$

The ability of the lambda calculus to “simulate recursion” in this way is the key to its power, and accounts for its persistence as a useful model of computation. Church recognized this, and is perhaps best expressed in his now famous thesis:

**Church’s Thesis:** *Effectively computable functions from positive integers to positive integers are just those definable in the lambda calculus.*

This is quite a strong claim! Although the notion of “functions from positive integers to positive integers” can be formalized very precisely, the notion of “effectively computable” cannot; thus no proof can be given for the thesis. However, it gained support from Kleene who in 1936 [Kle36] showed that “ $\lambda$ -definability” was precisely equivalent to Gödel and Herbrand’s notions of “recursiveness.” Meanwhile, Turing had been working on his now famous *Turing Machine* [Tur36], and in 1937 [Tur37] he showed that “Turing computability” was also precisely equivalent to  $\lambda$ -definability. These were quite satisfying results.<sup>9</sup>

The lambda calculus and the Turing machine were to have profound impacts on programming languages and computational complexity,<sup>10</sup> respectively, and computer science in general. This influence was probably much greater than Church or Turing could have imagined, which is perhaps not surprising given that computers did not even exist yet!

In parallel with the development of the lambda calculus, Schönfinkel and Curry were busy founding *combinatory logic*. It was Schönfinkel [Sch24] who discovered the surprising result that any function could be expressed as the composition of only two simple functions, K and S. Curry proved the consistency of a pure combinatory calculus [Cur30], and with Feys elaborated the theory considerably [CF58]. Although this work deserves as much attention from a logician’s point of view as the lambda calculus, and in fact its origins predate that of the lambda calculus, we will not pursue it further here, since it did not contribute directly to the development of functional languages in the way that the lambda calculus did. On the other hand, the combinatory calculus was to eventually play a surprising role in the *implementation* of functional languages, beginning with [Tur79] and summarized in [PJ87, Chapter 16].

Another noteworthy attribute of the lambda calculus is its restriction to functions of one argument. That it suffices to consider only such functions was first suggested by Frege in 1893 [vH67] and independently by Schönfinkel in 1924 [Sch24]. This restriction was later exploited by Curry [CF58], who used the notation  $(f x y)$  to denote  $((f x) y)$ , and which previously would have been written  $f(x, y)$ . This notation has become known as “currying,” and  $f$  is said to be a “curried function.” As we will see, the

---

<sup>9</sup>Much later Post [Pos43] and Markov [Mar51] proposed two other formal notions of effective computability, but these also were shown to be equivalent to  $\lambda$ -definability.

<sup>10</sup>Although the lambda calculus and the notion of  $\lambda$ -definability pre-dated the Turing Machine, complexity theorists latched onto the Turing Machine as their fundamental measure of decidability. This is probably because of the appeal of the Turing Machine as a *machine*, giving it more credibility in the emerging arena of electronic digital computers. See [Tra88] for an interesting discussion of this issue.

notion of currying has carried over today as a distinguishing syntactic characteristic of modern functional languages.

There are several variations and embellishments of the lambda calculus that historically could be mentioned here, but instead of doing that I will postpone their discussion until the point at which functional languages exhibited similar characteristics. In this way we can see more clearly the relationship between the lambda calculus and functional languages.

## 2.2 Lisp

A discussion of the history of functional languages would certainly be remiss if it did not include a discussion of Lisp, beginning with McCarthy's seminal work in the late 1950's.

Although lambda calculus is often considered as the foundation of Lisp, by McCarthy's own account [McC78] the lambda calculus actually played a rather small role. Its main impact came through McCarthy's desire to represent functions "anonymously," and Church's  $\lambda$ -notation was what he chose: a lambda abstraction written  $\lambda x.e$  in lambda calculus would be written `(lambda (x) e)` in Lisp.

Beyond that, the similarity wanes. For example, rather than use the *Y* combinator to express recursion, McCarthy invented the *conditional expression*,<sup>11</sup> with which recursive functions could be defined explicitly (and, arguably, more intuitively). As an example, the non-recursive factorial function given in the lambda calculus in the last section would be written recursively in Lisp in the following way:

```
(define fac (n)
  (if (= n 0)
      1
      (* n (fac (- n 1)))))
```

This and other ideas were described in two landmark papers in the early 60's [McC60, McC63], papers that inspired work on Lisp for many years to come.

McCarthy's original motivation for developing Lisp was the desire for an algebraic list processing language for use in artificial intelligence research. Although symbolic processing was a fairly radical idea at the time, his aims were quite pragmatic. One of the earliest attempts at designing such a language was suggested by McCarthy, and resulted in FLPL (for Fortran-compiled List Processing Language), implemented in 1958 on top of the Fortran system on the IBM 704 [GHG60]. Over the next few years McCarthy designed, refined, and implemented Lisp. His chief contributions during this period were:

---

<sup>11</sup>The conditional in Fortran (essentially the only other programming language in existence at the time) was a *statement*, not an expression, and was for *control*, not value-defining, purposes.

1. The conditional expression and its use in writing recursive functions.
2. The use of lists and higher-order operations over lists such as `mapcar`.
3. The central idea of a “cons cell,” and the use of garbage collection as a method of reclaiming unused cells.
4. The use of S-expressions (and abstract syntax in general) to represent both program and data.<sup>12</sup>

All four of these features are essential ingredients of any Lisp implementation today, and the first three are essential to functional language implementations as well.

A simple example of a typical Lisp definition is the following one for `mapcar`:

```
(define mapcar (fun lst)
  (if (null lst)
      nil
      (cons (fun (car lst)) (mapcar fun (cdr lst))) ))
```

This example demonstrates all of the points mentioned above. Note that the function `fun` is passed as an argument to `mapcar`. Although such higher-order programming was very well known in lambda calculus circles, it was certainly a radical departure from Fortran, and has become one of the most important programming techniques in Lisp and functional programming (higher-order functions are discussed more in Section 3.1). The primitive functions `cons`, `car`, `cdr`, and `null` are the well-known operations on lists whose names are still used today. `cons` creates a new list cell without burdening the user with explicit storage management; similarly, once that cell is no longer needed a “garbage collector” will come along and reclaim it, again without user involvement. For example, since `mapcar` constructs a new list from an old one, in the call:

```
(mapcar fun (cons a (cons b nil)))
```

the list `(cons a (cons b nil))` will become garbage after the call and will automatically be reclaimed. Lists were to become the paradigmatic data structure in Lisp and early functional languages.

The definition of `mapcar` in a modern functional language such as Haskell would appear similarly, except that pattern-matching would be used to “destructure” the list:

```
mapcar fun [] = []
mapcar fun (x:xs) = fun x : mapcar fun xs
```

---

<sup>12</sup>Interestingly, McCarthy claims that it was the read and print routines that influenced this notation most! [McC78]

[ [] is the null list, and : is the infix operator for cons; also note that function application has higher precedence than any infix operator.]

McCarthy was also interested in designing a practical language, and thus Lisp had many pragmatic features—in particular, sequencing, the assignment statement, and other primitives that induced side-effects on the store. Their presence undoubtedly had much to do with early experience with Fortran. Nevertheless, McCarthy emphasized the “mathematical elegance” of Lisp in his early papers, and in a much later paper his student Cartwright demonstrated the ease with which one could prove properties about pure Lisp programs [Car76].

Despite its impurities, Lisp had a great influence on functional language development, and it is encouraging to note that modern Lisps (especially Scheme) have returned more to the purity of the lambda calculus rather than the *ad hocery* that plagued the Maclisp era. This return to purity includes the “first-class” treatment of functions and the lexical scoping of identifiers. Furthermore, the preferred modern style of Lisp programming, such as espoused by Abelson and Sussman [ASS85], can be characterized as being predominantly side-effect free. And finally, we point out that Henderson’s Lispkit Lisp [Hen80] is a purely functional version of Lisp that uses an infix, algebraic syntax.

### 2.2.1 Lisp in Retrospect

Before continuing our historical development it is helpful to consider some of the design decisions McCarthy made and how they would be formalized in terms of the lambda calculus. It may seem that conditional expressions, for example, are an obvious feature to have in a language, but I think that only reflects our familiarity with modern high-level programming languages, most of which have them. In fact the lambda calculus version of the factorial example given in the previous section used a conditional (not to mention arithmetic operators) yet most readers probably understood it perfectly, and did not object to my departure from precise lambda calculus syntax!

The effect of conditional expressions can in fact be achieved in the lambda calculus by encoding the true and false values as functions, as well as defining a function to emulate the conditional:

$$\begin{aligned} \text{true} &\equiv \lambda x.\lambda y.x \\ \text{false} &\equiv \lambda x.\lambda y.y \\ \text{cond} &\equiv \lambda p.\lambda c.\lambda a.(p\ c\ a) \end{aligned}$$

In other words,  $(\text{cond } p\ c\ a) \equiv (\text{if } p\ \text{then } c\ \text{else } a)$ . One can then define, for example, the factorial function by:

$$\text{fac} \equiv \lambda n.\text{cond} (= n\ 0)\ 1\ (*\ n\ (\text{fac}\ (-\ n\ 1)))$$

where `=` is defined by:

$$\begin{aligned} (= n n) &\Rightarrow \text{true} \\ (= n m) &\Rightarrow \text{false, if } m \neq n \end{aligned}$$

where  $m$  and  $n$  range over the set of integer constants. However, I am still cheating a bit by not explaining the nature of the objects `-`, `*`, `0`, `1`, etc. in pure lambda calculus terms. It turns out that they can be represented in a variety of ways, essentially using functions to simulate the proper behavior, just as for `true`, `false`, and the conditional (for the details, see [Chu41]). In fact *any* conventional data or control structure can be “simulated” in the lambda calculus; if this were not the case, it would be difficult to believe Church’s Thesis!

However, even if McCarthy knew of these ways to express things in the lambda calculus (there is reason to believe that he did not), efficiency concerns might have rapidly led him to consider other alternatives anyway, especially since Fortran was the only high-level programming language anyone had any experience with. In particular, Fortran functions evaluated their arguments *before* entering the body of the function, resulting in what is often called a *strict*, or *call-by-value*, evaluation policy, corresponding roughly to applicative-order reduction in the lambda calculus. With this strategy extended to the primitives, including the conditional, one cannot easily define recursive functions. For example, in the above definition of factorial all three arguments to *cond* would be evaluated, including *fac* `(- n 1)`, resulting in non-termination.

*Non-strict* evaluation, corresponding to the normal-order reduction that is essential to the lambda calculus in realizing recursion, was not very well understood at the time—it was not at all clear how to implement it efficiently on a conventional von Neumann computer—and we would have to wait another 20 years or so before such an implementation was even attempted.<sup>13</sup> The conditional expression essentially allowed one to *selectively* invoke normal-order, or non-strict, evaluation. Stated another way, McCarthy’s conditional, although an expression, was compiled into code that essentially controlled the “reduction” process. Most imperative programming languages today that allow recursion do just that, and thus even though such languages are often referred to as strict, they all rely critically on at least one non-strict construct: the conditional.

### 2.2.2 The Lambda Calculus With Constants

The conditional expression is actually only one example of very many primitive functions that were included in Lisp. Rather than explain them in terms of the lambda calculus by a suitable encoding (i.e. compilation), it is perhaps better to formally extend the lambda calculus by adding a set of *constants* along with a set of what are usually

---

<sup>13</sup>On the other hand, the *call-by-name* evaluation strategy invented in Algol had very much of a normal-order reduction flavor. See [Weg68, Wad71] for early discussions of these issues.



called  $\delta$ -rules which state relationships between constants and which effectively extend the basis set of  $\alpha$ -,  $\beta$ -, and  $\eta$ -reduction rules. For example, the reduction rules for = that I gave earlier (and which are repeated below) are  $\delta$ -rules. This new calculus, often called the *lambda calculus with constants*, can be given a precise abstract syntax:

$x \in Id$	Identifiers
$c \in Con$	Constants
$e \in Exp$	Lambda Expressions

where  $e ::= x \mid c \mid e_1 e_2 \mid \lambda x.e$

for which various  $\delta$ -rules apply, such as the following:

$(= 0 0)$	$\Rightarrow$	$True$
$(= 0 1)$	$\Rightarrow$	$False$
...		
$(+ 0 0)$	$\Rightarrow$	$0$
$(+ 0 1)$	$\Rightarrow$	$1$
...		
$(+ 27 32)$	$\Rightarrow$	$59$
...		
$(If True e_1 e_2)$	$\Rightarrow$	$e_1$
$(If False e_1 e_2)$	$\Rightarrow$	$e_2$
...		
$(Car (Cons e_1 e_2))$	$\Rightarrow$	$e_1$
$(Cdr (Cons e_1 e_2))$	$\Rightarrow$	$e_2$
...		

where =, +, 0, 1, *If*, *True*, *False*, *Cons*, etc. are elements of *Con*.

The above rules can be shown to be a *conservative extension* of the lambda calculus, a technical term that in our context essentially means that convertible terms in the original system are still convertible in the new, and (perhaps more importantly) inconvertible terms in the original system are still inconvertible in the new. In general, care must be taken when introducing  $\delta$ -rules, since all kinds of inconsistencies could arise. For a quick and dirty example of inconsistency, we might define a primitive function over integers called *broken* with the following  $\delta$ -rules:

$(broken 0)$	$\Rightarrow$	$0$
$(broken 0)$	$\Rightarrow$	$1$

which immediately implies that more than one normal form exists for some terms, violating the first Church-Rosser theorem (see Section 2.1)!

As a more subtle example, suppose we define the logical relation *Or* by:

$(Or True e)$	$\Rightarrow$	$True$
$(Or e True)$	$\Rightarrow$	$True$
$(Or False False)$	$\Rightarrow$	$False$

Even though these rules form a conservative extension of the lambda calculus, a commitment to evaluate either of the arguments may lead to non-termination, even though the other argument may reduce to *True*. In fact it can be shown that with the above rules there does not exist a deterministic sequential reduction strategy that will guarantee that the normal form *True* will be found for all terms having such a normal form, and thus the second Church-Rosser property is violated. This version of *Or* is often called the “parallel or,” since a parallel reduction strategy is needed to implement it properly (and with which the first Church-Rosser Theorem will at least hold). Alternatively, we could define a “sequential or” by:

$$\begin{aligned}(Or\ True\ e) &\Rightarrow\ True \\(Or\ False\ e) &\Rightarrow\ e\end{aligned}$$

which can be shown to satisfy both Church-Rosser theorems.

### 2.3 Iswim

Historically speaking, Peter Landin’s work in the mid 60’s was the next significant impetus to the functional programming paradigm. Landin’s work was deeply influenced by that of Curry and Church. His early papers discussed the relationship between lambda calculus and both machines and high-level languages (specifically Algol 60). In [Lan64] Landin discussed how one could mechanize the evaluation of expressions, through an abstract machine called the SECD machine; and in [Lan65] he formally defined a non-trivial subset of Algol 60 in terms of the lambda calculus.

It is apparent from his work that Landin regarded highly the expressiveness and purity of the lambda calculus, and at the same time recognized its austerity. Undoubtedly as a result of this work, in 1966 Landin introduced a language (actually a family of languages) called Iswim (for “If you See What I Mean”), which included a number of significant syntactic and semantics ideas [Lan66]. Iswim, according to Landin, “can be looked on as an attempt to deliver Lisp from its eponymous commitment to lists, its reputation for hand-to-mouth storage allocation, the hardware dependent flavor of its pedagogy, its heavy bracketing, and its compromises with tradition.” When all is said and done, I think the primary contributions of Iswim, with respect to the development of functional languages, are the following:

1. Syntactic innovations:
  - (a) The abandonment of prefix syntax in favor of infix.
  - (b) The introduction of *let* and *where* clauses, including a notion of simultaneous and mutually recursive definitions.
  - (c) The use of an “off-side rule” based on *indentation* rather than separators (such as commas or semicolons) to scope declarations and expressions. For example (using Haskell syntax), the program fragment:



here: We assume that a block of declarations  $d_1 \cdots d_n$  always is potentially mutually recursive—if it isn't, our rules still work:

$$\begin{aligned}
 (\text{let } d_1 \cdots d_n \text{ in } e) &\Rightarrow (e \text{ where } d_1 \cdots d_n) \\
 (x \ x_1 \cdots x_n = e) &\Rightarrow (x = \lambda x_1. \lambda x_2. \cdots \lambda x_n. e) \\
 (e \text{ where } x_1 = e_1) &\Rightarrow (\lambda x_1. e)(Y \lambda x_1. e_1) \\
 (e \text{ where } (x_1 = e_1) \cdots (x_n = e_n)) &\Rightarrow (\lambda x_1. e)(Y \lambda x_1. e_1) \text{ where } x_2 = (\lambda x_1. e_2)(Y \lambda x_1. e_1) \\
 &\quad \cdots \\
 &\quad x_n = (\lambda x_1. e_n)(Y \lambda x_1. e_1)
 \end{aligned}$$

These rules are semantically equivalent to Landin's, but they avoid the need for a tupling operator to handle mutual recursion, and they use the Y combinator (defined in Section 2.1) instead of an iterative “unfolding” step.

We will call this resulting system the *recursive lambda calculus with constants*, or just *recursive lambda calculus*.

Landin's emphasis on expressing *what* the desired result is, as opposed to saying *how* to get it, and his claim that Iswim's declarative<sup>14</sup> style of programming was better than the incremental and sequential imperative style, were ideas to be echoed by functional programming advocates to this day. On the other hand, it took another ten years before interest in functional languages was to be substantially renewed. One of the reasons is that there were no decent implementations of Iswim-like languages around; this reason, in fact, was to persist into the 80's.

## 2.4 APL

Iverson's APL [Ive62], although not a purely functional programming language, is worth mentioning because its functional subset is an example of how one could achieve functional programming without relying on lambda expressions. In fact Iverson's design of APL was motivated out of his desire to develop an *algebraic* programming language for arrays, and his original work used an essentially functional notation. Subsequent development of APL resulted in several imperative features, but the underlying principles should not be overlooked.

APL was also unique in its goal of succinctness, in that it used a specially designed alphabet to represent programs—each letter corresponding to one operator. That APL became popular is apparent in the fact that many keyboards, both for typewriters and computer terminals, carried the APL alphabet. Backus' FP, which came after APL, was certainly influenced by the APL philosophy, and its abbreviated publication form also used a specialized alphabet (see the example in Section 2.5). In fact FP has much in common with APL, the primary difference being that FP's fundamental data structure is the *sequence*, whereas in APL it is the *array*.

<sup>14</sup>Landin actually disliked the term “declarative,” preferring instead “denotative.”

It is worth noting that recent work on APL has revived some of APL's purely functional foundations. The most notable work is that of Tu [Tu86, TP86], who designed a language called FAC, for Functional Array Calculator (presumably a take-off on Turner's KRC, Kent Recursive Calculator). FAC is a purely functional language that adopts most of the APL syntax and programming principles, but in addition has special features to allow programming with *infinite* arrays; naturally, lazy evaluation plays a major role in the semantics of the language. Another interesting approach is that of Mullin [Mul88].

## 2.5 FP

Backus' FP was one of the earliest functional languages to receive wide-spread attention. Although most of the features in FP are not found in today's modern functional languages, Backus' Turing Award lecture [Bac78] in 1978 was one of the most influential and now most-often cited papers extolling the functional programming paradigm. It not only stated quite eloquently why functional programming was "good," but also quite vehemently why traditional imperative programming was "bad."<sup>15</sup> Backus' coined the term "word-at-a-time programming" to capture the essence of imperative languages, showed how such languages were inextricably tied to the "von Neumann machine," and gave a convincing argument why such languages were not going to meet the demands of modern software development. That this argument was being made by the person who is given the most credit for designing Fortran, and who also had significant influence on the development of Algol, led substantial weight to the "functional thesis." The exposure given to Backus' paper was one of the best things that could have happened to the field of functional programming, which at the time was certainly not considered main-stream.

Despite the great impetus Backus' paper gave to functional programming, it is interesting to note that in the same paper Backus also said that languages based on lambda calculus would have problems, both in implementation and expressiveness, because the model was not suitably history-sensitive (more specifically, it did not handle large data structures, such as databases, very easily). With regard to implementation this argument is certainly understandable, because it was not clear how to implement the notion of substitution in an efficient manner, nor was it clear how to structure data in such ways that large data structures could be implemented efficiently (both of these issues are much better understood today). And with regard to expressiveness, well, that argument is still a matter of debate today. In any case, these problems were the motivation for Backus' Applicative State Transition (AST) Systems, in which state is introduced as something on which purely functional programs interact with in a more traditional (i.e. imperative) way.

Perhaps more surprising, and an aspect of the paper that is usually overlooked, Backus had this to say about "lambda-calculus based systems:"

---

<sup>15</sup>Ironically, the Turing Award was given to Backus in a large part because of his work on Fortran!

“An FP system is founded on the use of a fixed set of combining forms called functional forms. . . . In contrast, a lambda-calculus based system is founded on the use of the lambda expression, with an associated set of substitution rules for variables, for building new functions. The lambda expression (with its substitution rules) is capable of defining all possible computable functions of all possible types and of any number of arguments. This freedom and power has its disadvantages as well as its obvious advantages. It is analogous to the power of unrestricted control statements in conventional languages: with unrestricted freedom comes chaos. If one constantly invents new combining forms to suit the occasion, as one can in the lambda calculus, one will not become familiar with the style or useful properties of the few combining forms that are adequate for all purposes.”

Backus’ argument, of course, was in the context of garnering support for FP, which had a small set of combining forms that were claimed to be sufficient for most programming applications. One of the advantages of this approach is that each of these combining forms could be named with particular brevity, and thus programs become quite compact—this was exactly the approach taken by Iverson in designing APL (see Section 2.4). For example, an FP program for inner product looks like:

$$\text{Def } IP \equiv (/ +) \circ (\alpha \times) \circ Trans$$

where  $/$ ,  $\circ$ , and  $\alpha$  are combining forms called “insert,” “compose,” and “apply-to-all,” respectively. In a modern functional language such as Haskell this would be written with slightly more verbosity as:

```
ip l1 l2 = foldl (+) 0 (map2 (*) l1 l2)
```

[In Haskell an infix operator such as  $+$  may be passed as an argument by surrounding it in parentheses.] Here `foldl` is the equivalent of insert ( $/$ ), and `map2` is a two-list equivalent of apply-to-all ( $\alpha$ ), thus eliminating the need for *Trans*. These functions are pre-defined in Haskell, as they are in most modern functional languages, for the same reason that Backus argues—they are commonly used. If they were not they could easily be defined by the user. For example, we may wish to define an infix composition operator for functions, the first of which is binary, as follows:

$$(f \text{ .o. } g) \times y = f (g \times y)$$

[Note how infix operators may be defined in Haskell; operators are distinguished lexically by their first character being non-alphabetic.] With this we can reclaim much of FP’s succinctness in defining `ip`:

```
ip = foldl (+) 0 .o. map2 (*)
```

[Recall that in Haskell function application has higher precedence than infix operator application.] It is for this reason, together with the fact that FP's specialization precluded the generality afforded by user-defined higher-order functions (which is all that combining forms are), that modern functional languages did not follow the FP style. As we shall soon see, certain other kinds of syntactic sugar became more popular instead (such as pattern-matching, list comprehensions, and sections).

Many extensions to FP have been proposed over the years, including the inclusion of strong typing and abstract datatypes [GHW81]. In much more recent work, Backus, Williams and Wimmers have designed the language FL [BWW86], which is strongly (although dynamically) typed, and in which higher-order functions and user-defined datatypes are allowed. Its authors still emphasize the algebraic style of reasoning that is paramount in FP, although it is also given a denotational semantics which is provably consistent with respect to the algebraic semantics.

## 2.6 ML

In the mid 70's, at the same time Backus was working on FP at IBM, several research projects were underway in the United Kingdom that related to functional programming, most notably work at Edinburgh. There Gordon, Milner, and Wadsworth had been working on a proof generating system called LCF for reasoning about recursive functions, in particular in the context of programming languages [GMW79]. The system consisted of a deductive calculus called  $PP\lambda$  (polymorphic predicate calculus) together with an interactive programming language called ML, for "meta-language" (since it served as the command language for LCF).

LCF is quite interesting as a proof system, but its authors soon found that ML was also interesting in its own right, and they proceeded to develop it as a stand-alone functional programming language [GMM\*78]. That it was, and still is, called a *functional* language is actually somewhat misleading, since it has a notion of *references* which are locations that can be stored into and read from, much as variables are assigned and read. Its I/O system also induces side-effects, and is not referentially transparent. Nevertheless, the style of programming that it encourages is still functional, and that is the way it was promoted (the same is somewhat true for Scheme, although to a lesser extent).

More recently a standardization effort for ML has taken place, in fact taking some of the good ideas of Hope [BMS80] along with it (such as pattern-matching), yielding a language now being called *Standard ML*, or just SML [Mil84, Wik88].

ML is a fairly complete language—certainly the most practical functional language at the time it appeared—and SML is even more so. It has higher-order functions, a simple I/O facility, a sophisticated module system, and even exceptions. But by far the most significant aspect of ML is its *type system*, which is manifested in several ways:

1. It is *strongly* and *statically* typed.
2. It uses *type inference* to determine the type of every expression, instead of relying on explicit type declarations.
3. It allows *polymorphic* functions and data structures; that is, functions may take arguments of arbitrary type if in fact the function does not depend on that type (similarly for data structures).
4. It has user-defined concrete and abstract datatypes (an idea actually borrowed from Hope, and not present in the initial design of ML).

ML was the first language to utilize type inference as a semantically integrated component of the language, and at the same time its type system was richer than any previous statically typed language in that it permitted true polymorphism. It seemed that the best of two worlds had been achieved—not only is making explicit type declarations (a sometimes burdensome task) not required, but in addition a program that successfully passes the type inferencer is guaranteed not to have any type errors. Unfortunately, this idyllic picture is not completely accurate (although it is close!), as we shall soon see.

We shall next discuss the issue of types in a foundational sense, thus continuing our plan of elaborating the lambda calculus to match the language features being discussed. This will lead us to a reasonable picture of ML’s *Hindley-Milner type system*, the rich polymorphic type system that was mentioned above, and that was later adopted by every other statically-typed functional language, including Miranda and Haskell. Aside from the type system, the two most novel features in ML are its *references* and *modules*, which are also covered in this section. Discussion of ML’s data abstraction facilities will be postponed until Section 3.3.

### 2.6.1 The Hindley-Milner Type System

We can introduce types into the pure lambda calculus by first introducing a domain of basic types, say *BasTyp*, as well as a domain of derived types, say *Typ*, and then requiring that every expression be “tagged” with a member of *Typ*, which we do by superscripting, as in “ $e^\tau$ .” The derived type “ $\tau_2 \rightarrow \tau_1$ ” denotes the type of all functions from values of type  $\tau_2$  to values of type  $\tau_1$ , and thus a proper application will have the form “ $(e_1^{\tau_2 \rightarrow \tau_1} e_2^{\tau_2})^{\tau_1}$ .” Modifying the pure lambda calculus in this way, we arrive at the



*pure typed lambda calculus:*

$b \in \text{BasTyp}$	Basic types
$\tau \in \text{Typ}$	Derived types
where $\tau ::= b \mid \tau_1 \rightarrow \tau_2$	
$x^\tau \in \text{Id}$	Typed identifiers
$e \in \text{Exp}$	Typed lambda expressions
where $e ::= x^\tau$	
$\mid (e_1^{\tau_2 \rightarrow \tau_1} e_2^{\tau_2})^{\tau_1}$	
$\mid (\lambda x^{\tau_2}. e^{\tau_1})^{\tau_2 \rightarrow \tau_1}$	

for which we then provide the following reduction rules:

1. Typed- $\alpha$ -conversion:  $(\lambda x_1^{\tau_1}. e^\tau) \Leftrightarrow (\lambda x_2^{\tau_1}. [x_2^{\tau_1}/x_1^{\tau_1}]e^\tau)$ , where  $x_2^{\tau_1} \notin \text{fv}(e^\tau)$ .
2. Typed- $\beta$ -conversion:  $((\lambda x^{\tau_2}. e_1^{\tau_1}) e_2^{\tau_2}) \Leftrightarrow [e_2^{\tau_2}/x^{\tau_2}]e_1^{\tau_1}$ .
3. Typed- $\eta$ -conversion:  $\lambda x^{\tau_1}. (e^{\tau_2} x^{\tau_1}) \Leftrightarrow e^{\tau_2}$ , if  $x^{\tau_1} \notin \text{fv}(e^{\tau_2})$ .

To preserve type correctness we assume that the typed identifiers  $x^{\tau_1}$  and  $x^{\tau_2}$ , where  $\tau_1 \neq \tau_2$ , are *distinct identifiers*. Note then how every expression in our new calculus carries with it its proper type, and thus type-checking is “built-in” to the syntax.

Unfortunately, there is a serious problem with this result: our new calculus has lost the power of  $\lambda$ -definability! In fact, every term in the pure typed lambda calculus can be shown to be *strongly normalizable*, meaning each has a normal form, and there is an effective procedure for reducing each of them to its normal form [FLO85]. Consequently we can only compute a rather small subset of functions from integers to integers—namely the extended polynomials [Bar84].<sup>16</sup>

The reader can gain some insight into the problem by trying to write the definition of the  $Y$  combinator—that paradoxical entity that allowed us to express recursion—as a typed term. The difficulty lies in properly typing self-application (recall the discussion in Section 2.1), since typing “ $(e e)$ ” requires that  $e$  have both the type  $\tau_2 \rightarrow \tau_1$  and  $\tau_2$ , which cannot be done within the structure we have given. It can, in fact, be shown that the pure typed lambda calculus has *no* fixpoint operator.

Fortunately, there is a clever way to solve this dilemma. Instead of relying on self application to implement recursion, simply add to the calculus a family of constant fixpoint operators similar to  $Y$ , only typed. To do this we first move into the *typed lambda calculus with constants*, in which a domain of constants  $Con$  is added as in

<sup>16</sup>On the bright side, some researchers view the strong normalization property as a *feature*, since it means that all programs are guaranteed to terminate. Indeed this property forms the basis of much of the recent work on using “constructive type theory” as a foundation for programming languages.

the (untyped) lambda calculus with constants. We then include in *Con* a *typed fixpoint operator* of the form  $Y_\tau$  for every type  $\tau$ , where:

$$Y_\tau : (\tau \rightarrow \tau) \rightarrow \tau$$

Then for each fixpoint operator  $Y_\tau$  we add the  $\delta$ -rule:

$$\text{Typed-}Y\text{-conversion: } (Y_\tau e^{\tau \rightarrow \tau})^\tau \Leftrightarrow (e^{\tau \rightarrow \tau} (Y_\tau e^{\tau \rightarrow \tau})^\tau)^\tau$$

The reader may easily verify that type consistency is maintained by this rule.

By ignoring the type information, we can see that the above  $\delta$ -rule corresponds to the conversion  $(Y f) \Leftrightarrow (f (Y f))$  in the untyped case, and the same trick for implementing recursion with  $Y$  as discussed in Section 2.1 can be used here, thus regaining  $\lambda$ -definability. For example, a non-recursive definition of the *typed* factorial function would be the same as the untyped version given earlier, except that  $Y_\tau$  would be used where  $\tau = \text{Int} \rightarrow \text{Int}$ .

In addition to this calculus, we can derive a *typed recursive lambda calculus with constants* in the same way that we did in the untyped case, except that instead of the unrestricted  $Y$  combinator we use the typed versions defined above.

At this point our type system is about at par with that of a strongly and explicitly typed language such as Pascal or Ada. However, we would like to do better. As mentioned in Section 2.6, the designers of ML extended the state-of-the-art of type systems in two significant ways:

- They permitted *polymorphic* functions.
- They used *type inference* to infer types rather than requiring that they be declared explicitly.

As an example of a polymorphic function, consider:

```
map :: (a -> b) -> [a] -> [b]
map f []      = []
map f (x:xs) = f x : map f xs
```

The first line is a *type signature* that declares the type of `map`; it can be read as “for all types `a` and `b`, `map` is a function that takes two arguments, a function from `a` into `b` and a list of elements of type `a`, and returns a list of elements of type `b`.”

[Type signatures are optional in Haskell; the type system is able to infer them automatically. Also note that the type constructor `->` is right associative, which is consistent with the left associativity of function application.]

So `map` can be used to map square down a list of integers, or head down a list of lists, etc. In a monomorphic language such as Pascal one would have to define a separate function for each of these. The advantage of polymorphism should be clear.

One way to achieve polymorphism in our calculus is to add a domain of *type variables* and extend the domain of derived types accordingly:

$$\begin{array}{ll}
 b \in \text{BasTyp} & \text{Basic types} \\
 v \in \text{TypId} & \text{Type variables} \\
 \tau \in \text{Typ} & \text{Derived types} \\
 \text{where } \tau ::= b \mid v \mid \tau_1 \rightarrow \tau_2 &
 \end{array}$$

Thus, for example,  $v \rightarrow \tau$  is a proper type, and can be read as if the type variable was universally quantified: “for all types  $v$ , the type  $v \rightarrow \tau$ .” To accommodate this we must change the rule for  $\beta$ -conversion to read:

2. Typed- $\beta$ -conversion with type variables:

$$\begin{array}{l}
 \text{(a) } ((\lambda x^{\tau_2}. e_1^{\tau_1}) e_2^{\tau_2}) \Leftrightarrow ([e_2^{\tau_2}/x^{\tau_2}]e_1^{\tau_1})^{\tau_1} \\
 \text{(b) } ((\lambda x^v. e_1^{\tau_1}) e_2^{\tau_2}) \Leftrightarrow ([\tau_2/v]([e_2^{\tau_2}/x^v]e_1^{\tau_1}))^{\tau_1}
 \end{array}$$

where substitution on type variables is defined in the obvious way. Note that this rule implies the validity of expressions of the form  $(e_1^{v \rightarrow \tau_1} e_2^{\tau_2})^{\tau_1}$ . Similar changes are required to accommodate expressions such as  $(e_1^{\tau \rightarrow v} e_2^{\tau})^v$ .

But alas, now that type variables have been introduced, it is no longer clear whether a program is properly typed—it is not “built-in” to the static syntactic structure of the calculus. In fact, the type checking problem for this calculus is undecidable, being a variant of a problem known as *partial* polymorphic type inference [Boe85, Pfe88].

Rather than trying to deal with the type checking problem directly, we might go one step further with our calculus and try to attain ML’s lack of a requirement for explicit typing. We can achieve this by simply *erasing all type annotations* and then trying to solve the seemingly harder problem of inferring types of completely “naked” terms. Surprisingly, it is not known whether an effective type inference algorithm exists for this calculus, even though the problem of partial polymorphic type inference, known to be undecidable, seems as if it should be easier.

Fortunately, Hindley [Hin69] and Milner [Mil78] independently discovered a restricted polymorphic type system that is almost as rich as that provided by our calculus, and for which type inference is decidable. In other words, there exist certain terms in the calculus presented above that one could argue are properly typed, but would not be allowed in the Hindley-Milner system. The system still allows polymorphism, such as exhibited by `map` defined earlier, and is able to infer the type of functions such as `map` without any type signatures present. However, the *use* of such polymorphism is limited to the scope in which `map` was defined. For example, the following program:

```

silly map f g
where f  :: Int -> Int
      g  :: Char -> Char
      map :: (a -> b) -> [a] -> [b]
      silly m f g = (m f num_list, m g char_list)

```

[“(e1, e2)” is a tuple] results in a type error, since `map` is *passed as an argument* and then instantiated in *two different ways*; i.e. once as type  $(\text{Int} \rightarrow \text{Int}) \rightarrow [\text{Int}] \rightarrow [\text{Int}]$  and once as type  $(\text{Char} \rightarrow \text{Char}) \rightarrow [\text{Char}] \rightarrow [\text{Char}]$ . If `map` were instantiated in several ways within the scope in which it was defined, or if `m` were only instantiated in one way within the function `silly`, there would have been no problem.

This example demonstrates a fundamental limitation to the Hindley-Milner type system, but in practice the class of programs that the system rejects is not large, and is certainly smaller than that rejected by any existing type-checking algorithm for conventional languages in that, if nothing else, it allows polymorphism. Many other functional languages have since then incorporated what amounts to a Hindley-Milner type system, including Miranda and Haskell. It is beyond the scope of this article to discuss the details of type inference, but the reader may find good pragmatic discussions in [Han] and [DM82], and a good theoretical discussion in [Mil78].

As a final comment we point out that an alternative way to gain polymorphism is to introduce types as *values* and give them at least some degree of “first-class status” (as we did earlier for functions), for example allowing them to be passed as arguments and returned as results. Allowing them to be passed as arguments only (and then used in type annotations in the standard way) results in what is known as the *polymorphic* or *second-order* typed lambda calculus. Girard [Gir72] and Reynolds [Rey74] discovered and studied this type system independently, and it appears to have great expressive power. However, it turns out to be essentially equivalent to the system we developed earlier, and has the same difficulties with respect to type inference. It is, nevertheless, an active area of current research (see [CW85] and [Rey85] for good summaries of this work).

## 2.6.2 ML’s References

A reference is essentially a pointer to a cell containing values of a particular type; they are created by the (pseudo-)function `ref`. For example, “`ref 5`” evaluates to an *integer reference*—a pointer to a cell that is allowed to contain only integers, and in this case having initial contents 5. The contents of a cell can be “read” using the prefix operator `!`. Thus if `x` is bound to “`ref 5`” then “`!x`” returns 5.

The cell pointed to by a reference may be updated via *assignment* using the infix operator `:=`. Continuing with the above example, “`x := 10`”, although an expression,

has the side-effect of updating the cell pointed to by `x` with the value 10. Subsequent evaluations of “!`x`” will then return 10. Of course, to make the notion of “subsequent” well-defined, it is necessary to introduce sequencing constructs; indeed ML even has an iterative `while` construct.

References in ML amount to assignable variables in a conventional programming language, and are only notable in that they can be included within a type structure such as Hindley-Milner’s, and can be relegated to a minor role in a language that is generally proclaimed as being functional. A proper treatment of references within a Hindley-Milner type system can be found in [Tof88].

### 2.6.3 Modules

Modules in ML are called *structures*, and are essentially reified *environments*. The type of a structure is captured in its *signature*, and contains all of the static properties of a module that are needed by some other module that might use it. The use of one module by another is captured by special functions called *functors* that map structures to new structures. This capability is sometimes called a “parameterized module” or “generic package.”

For example, a new signature called SIG (think of it as a new type) may be declared by:

```
signature SIG =
  sig
    val x : int
    val succ : int -> int
  end
```

in which the types of two identifiers have been declared, `x` of type `int` and `succ` of type `int->int`. The following structure `S` (think of it as an environment) has the implied signature SIG defined above:

```
structure S =
  struct
    val x = 5
    val succ x = x+1
  end
```

If we then define the following functor `F` (think of it as a function from structures to structures):

```

functor F(T:SIG) =
  struct
    val y = T.x + 1
    val add2 x = T.succ(T.succ(x))
  end

```

then the new structure declaration:

```

structure U = F(S)

```

is equivalent to having written:

```

structure U =
  struct
    val y = x + 1
    val add2 x = succ(succ(x))
    val x = 5
    val succ x = x+1
  end

```

except that the signature of U does not include bindings for x and succ (i.e. they are hidden).

Although seemingly simple, the ML module facility has one very noteworthy feature: Structures are (at least partially) “first-class” in that functors take them as arguments and return them as values. A more conservative design (such as adopted in Miranda and Haskell) might require all modules to be *named*, thus relegating them to second-class status. Of course, this first-class treatment has to be ended somewhere if type-checking is to remain effective, and in ML that is manifested in the fact that structures can be passed to functors *only* (for example they cannot be placed in lists), the signature declaration of a functor’s argument is *mandatory* in the functor declaration, and *functors* themselves are not first-class.

It is not clear whether this “almost-first-class” treatment of structures is worth the extra complexity, but the ML module facility is certainly the most sophisticated of those found in existing functional programming languages, and it is achieved with no loss of static type-checking ability, including the fact that modules may be compiled independently and later “linked” via functor application.

## 2.7 SASL, KRC, and Miranda

At the same time ML and FP were being developed, David Turner, first at the University of St. Andrews and later at the University of Kent, was busy working on another

style of functional languages resulting in a series of three languages that characterize most faithfully the “modern school” of functional programming ideas. More than any other researcher, Turner argued eloquently for the value of lazy evaluation, higher-order functions, and the use of recursion equations as a syntactic sugaring for the lambda calculus [Tur81, Tur82]. Turner’s use of recurrence equations was consistent with Landin’s argument ten years earlier, as well as Burge’s excellent treatise on recursive programming techniques [Bur75] and Burstall and Darlington’s work on program transformation [BD77]. But the prodigious use of higher-order functions and lazy evaluation, especially the latter, was something new, and was to become a hallmark of modern functional programming techniques.

In the development of SASL (St. Andrews Static Language) [Tur76], KRC (Kent Recursive Calculator) [Tur81], and Miranda<sup>17</sup> [Tur85], Turner concentrated on making things easier on the programmer, and thus he introduced various sorts of syntactic sugar. In particular, using SASL’s syntax for equations gave programs a certain mathematical flavor, since equations were deemed applicable through the use of *guards* and Landin’s notion of the “off-side rule” was revived. For example, this definition of the factorial function:

$$\begin{aligned} \text{fac } n &= 1, && n=0 \\ &= n * \text{fac}(n-1), && n>0 \end{aligned}$$

looks a lot like the mathematical version:

$$\text{fac } n = \begin{cases} 1 & \text{if } n = 0 \\ n * \text{fac}(n - 1) & \text{if } n > 0 \end{cases}$$

(In Haskell this program would be written with slightly different syntax as:

$$\begin{aligned} \text{fac } n \mid n==0 &= 1 \\ ' &\mid n>0 = n * \text{fac}(n-1) \end{aligned}$$

More on equations and pattern-matching may be found in Section 3.4.)

Another nice aspect of SASL’s equational style is that it facilitates the use of higher-order functions through *currying*. For example, if we define:

$$\text{add } x \ y = x+y$$

then “add 1” is a function that adds one to its argument.

KRC is an embellishment of SASL primarily through the addition of *ZF expressions* (which were intended to resemble Zemelo-Frankel set abstraction and whose syntax was originally suggested by John Darlington), as well as various other short-hands for lists (such as [a..b] to denote the list of integers from a to b, and [a..] to denote the “infinite” sequence starting with a). For example (using Haskell syntax):

---

<sup>17</sup>Miranda is one of the few (perhaps the only) functional languages to be marketed commercially.

```
[ x*x | x <- [1..100], odd(x) ]
```

is the list of squares of the odd numbers from 1 to 100, and is similar to:

$$\{x^2 \mid x \in \{1, 2, \dots, 100\} \wedge \text{odd}(x)\}$$

except that the former is a list, the latter is a set. In fact Turner used the term “set abstraction” as synonymous with “ZF expression,” but in fact both terms are somewhat misleading since the expressions actually denote lists, not true sets. The more popular current term is “list comprehension,”<sup>18</sup> which is what I will call them in the remainder of this paper. As an example of the power of list comprehensions, here is a rather concise and perspicuous definition of *quicksort*:

```
qs []      = []
qs (x:xs) = qs [y | y<-xs, y<x] ++ [x] ++
             qs [y | y<-xs, y>=x]
```

[++ is the infix *append* operator.]

Miranda is in turn an embellishment of KRC, primarily in its treatment of types: it is strongly typed, using a Hindley-Milner type system, and it allows user defined concrete and abstract datatypes (both of these ideas were presumably borrowed from ML; see Sections 2.6 and 2.6.1). One interesting innovation in syntax in Miranda is its use of *sections* (first suggested by Richard Bird), which are a convenient way to convert partially applied infix operators into functional values. For example, the expressions (+), (x+), and (+x) correspond to the functions f, g, and h, respectively, defined by:

```
f x y = x+y
g y    = x+y
h y    = y+x
```

Partly because of the presence of sections, Miranda does not provide syntax for lambda abstractions. (In contrast, the Haskell designers chose to have lambda abstractions, and thus chose not to have sections.)

Turner was perhaps the foremost proponent of both higher-order functions and lazy evaluation, although the ideas originated elsewhere. As I have done with data abstraction, discussion of both of these topics is delayed until Sections 3.1 and 3.2, respectively, where they are discussed in a more general context.

---

<sup>18</sup>A term popularized by Phil Wadler.



## 2.8 Dataflow Languages

In the area of computer architecture, beginning predominantly with Jack Dennis' work in the early 70's [DM74], there arose the notion of *dataflow*, a computer architecture organized solely around the data dependencies in a program, resulting in high degrees of parallelism. Since data dependencies were paramount, and artificial sequentiality was objectionable, the languages designed to support such machines were essentially functional languages, although historically they have been called *dataflow languages*. In fact they do have a few distinguishing features, typically reflecting the idiosyncrasies of the dataflow architecture (just as imperative languages reflect the von Neumann architecture): they are typically first-order (reflecting the difficulty in constructing closures in the dataflow model), strict (reflecting the data-driven mode of operation that was most popular and easiest to implement), and in certain cases do not even allow recursion (reflecting Dennis' original *static* dataflow design, rather than, for example, Arvind's *dynamic* tagged-token model [AG77, AK81]). A good summary of work on dataflow machines, at least through 1981, can be found in [TBH82]; more recently, see [Veg84].

The two most important dataflow languages developed during this era were Dennis et al.'s Val [AD79, Mcg82] and Arvind et al.'s Id [AG82]. More recently, Val has evolved into SISAL [MAGD83], and Id into Id Nouveau [NPA86]. The former has retained much of the strict and first-order semantics of dataflow languages, whereas the latter has many of the features that characterize modern functional languages.

Keller's FGL [KJRL80] and Davis' DDN [Dav78] are also notable developments that accompanied a flurry of activity on dataflow machines at the University of Utah in the late 70's. Yet another interesting dataflow language is Ashcroft and Wadge's Lucid [AW76a, AW76b, WA85], whose distinguishing feature is the use of identifiers to represent streams of values (in a temporal, dataflow sense), thus allowing the expression of iteration in a rather concise manner. The authors also developed an algebra for reasoning about Lucid programs.

## 2.9 Others

In the late 70's and early 80's a surprising number of other modern functional languages appeared, most in the context of implementation efforts. These included Hope at Edinburgh University [BMS80], FEL at Utah [Kel82], Lazy ML (LML) at Chalmers [Aug84], Alfl at Yale [Hud84], Ponder at Cambridge [Fai85], Orwell at Oxford [WM88], Daisy at Indiana [Joh88], Twentel at the University of Twente [Kro87], and Tui at Victoria University [Bou88].

Perhaps the most notable of these languages was Hope, designed and implemented by Rod Burstall, David MacQueen, and Ron Sannella at Edinburgh University [BMS80].

Their goal was “to produce a very simple programming language which encourages the production of clear and manipulable programs.” Hope is strongly typed, allows polymorphism, but requires explicit type declarations as part of all function definitions (which also allowed a useful form of overloading). It has lazy lists, but otherwise is strict. It also has a simple module facility. But perhaps the most significant aspect of Hope is its user-defined concrete datatypes and the ability to pattern-match against them. ML, in fact, did not originally have these features; they were borrowed from Hope in the design of SML.

To quote Bird and Wadler [BW88], this proliferation of functional languages was “a testament to the vitality of the subject,” although by 1987 there existed so many functional languages that there truly was a “Tower of Babel,” and something had to be done! The funny thing was, the semantic underpinnings of these languages were fairly consistent, and thus the researchers in the field had very little trouble understanding each other’s programs, so the motivation within the *research* community to standardize on a language was not high.

Nevertheless, in September of 1987 a meeting was held at the FPCA Conference in Portland, Oregon, to discuss the problems that this proliferation of languages was creating. There was a strong consensus that the general use of modern, non-strict functional languages was being hampered by the lack of a common language. Thus it was decided that a committee should be formed to design such a language, providing faster communication of new ideas, a stable foundation for real applications development, and a vehicle through which other people would be encouraged to learn and use functional languages. The result of that committee’s effort was a purely functional programming language called Haskell [HWe88], named after Haskell B. Curry, and described in the next section.

## 2.10 Haskell

Haskell is a general purpose, purely functional programming language exhibiting many of the recent innovations in functional (as well as other) programming language research, including higher-order functions, lazy evaluation, static polymorphic typing, user-defined datatypes, pattern-matching, and list comprehensions. It is also a very *complete* language in that it has a module facility, a well-defined functional I/O system, and a rich set of primitive datatypes, including lists, arrays, arbitrary and fixed precision integers, and floating-point numbers. In this sense Haskell represents both the culmination and solidification of many years of research on functional languages—the design was influenced by languages as old as Iswim and as new as Miranda.

Haskell also has several interesting new features; most notably, a systematic treatment of *overloading*, an “orthogonal” abstract datatype facility, a universal and purely functional I/O system, and, by analogy to list comprehensions, a notion of *array comprehensions*.

Haskell is not a small language. The decision to emphasize certain features such as pattern-matching and user-defined datatypes, and the desire for a *complete* and *practical* language that includes such things as I/O and modules, necessitates a somewhat large design. The Haskell Report also provides a denotational semantics for both the static and dynamic behavior of the language; it is considerably more complex than the simple semantics defined in Section 3.5 for the lambda calculus, but then again one wouldn't really want to program in as sparse a language as the lambda calculus.

Will Haskell become a standard? Will it succeed as a useful programming language? Only time will tell. As with any other language development, it is not only the quality of the design that counts, but also the ready availability of good implementations, and the backing from vendors, government agencies, and researchers alike. At this date it is too early to tell what role each of these factors will play.

I will end our historical development of functional languages here, without elaborating on the details of Haskell just yet. Those details will surface in significant ways in the next section, where the most important features of modern function languages are discussed; and in the following section, where more advanced ideas and active research areas are discussed.

### 3 Distinguishing Features of “Modern” Functional Languages

The reader will recall that I chose to delay detailed discussion of four distinguishing features of modern functional languages—higher-order functions, lazy evaluation, data abstraction mechanisms, and equations/pattern-matching. Now that we have completed our study of the historical development of functional languages, we can return to those features. Most of the discussion will center on how the features are manifested in Haskell, ML, and Miranda.

#### 3.1 Higher-Order Functions

If functions are treated as “first-class values” in a language—allowing them to be stored in data structures, passed as arguments, and returned as results—they are referred to as *higher-order functions*. I have not said too much about the utility of higher-order functions thus far, although they exist in most of the functional languages that we've discussed, including of course the lambda calculus. Their utility has in fact been argued in many circles, including ones outside of functional programming, most notably the Scheme community (see, for example, [ASS85]).

The main philosophical argument for higher-order functions is that functions are values just like any others, so why not give them the same “first-class” status? But

there are also compelling pragmatic reasons for wanting higher-order functions. Simply stated, the function is the primary *abstraction mechanism over values*, and thus facilitating the use of functions increases the utility of that kind of abstraction.

As an example of a higher-order function, consider:

$$\text{twice } f \ x = f (f \ x)$$

which takes its first argument, a function  $f$ , and “applies it twice” to its second argument,  $x$ . The syntax used here is important: `twice` as written is *curried*, meaning that when applied to one argument it returns a function which then takes one more argument, the second argument above. For example, the function `add2`:

$$\begin{aligned} \text{add2} &= \text{twice succ} \\ &\text{where succ } x = x+1 \end{aligned}$$

is a function that will add 2 to its argument. Making function application associate to the left facilitates this mechanism, since “`twice succ x`” is equivalent to “(`twice succ`)  $x$ ” so everything works out just fine.

In modern functional languages functions can be created in several ways. One way is to name them using equations, as above, but another way is to create them directly as *lambda abstractions*, thus rendering them nameless, as in the Haskell expression:

$$\backslash x \rightarrow x+1$$

[in lambda calculus this would be written  $\lambda x.x + 1$ ] which is the same as the successor function `succ` defined above. `add2` can then be defined more succinctly as:

$$\text{add2} = \text{twice } (\backslash x \rightarrow x+1)$$

From a pragmatic viewpoint we can understand the utility of higher-order functions by analyzing the utility of abstraction in general. As we all know from introductory programming, a function is an abstraction of values over some common behavior (an expression). Limiting the values over which the abstraction occurs to non-functions seems unreasonable; lifting that restriction results in higher-order functions. Hughes makes a slightly different but equally compelling argument in [Hug84], where he emphasizes the importance of *modularity* in programming and argues convincingly that higher-order functions increase modularity by serving as a mechanism for “glueing” program fragments together. That glueing property comes not just from the ability to *compose* functions, but also the ability to *abstract* over functional behavior as described above.

As an example, suppose in the course of program construction we define a function to add together the elements of a list:

```
sum [] = 0
sum (x:xs) = add x (sum xs)
```

Then suppose we later define a function to multiply the elements of a list:

```
prod [] = 1
prod (x:xs) = mul x (prod xs)
```

But now we notice a repeating pattern, and anticipate that we might see it again, so we ask ourselves if we can possibly abstract the common behavior. In fact this is easy to do: we note that `add/mul` and `0/1` are the variable elements in the behavior, and thus we *parameterize* them; that is, we make them formal parameters, say `f` and `init`. Calling the new function `fold`, the equivalent of `sum/prod` will be `fold f init`, and thus we arrive at:

```
(fold f init) [] = init
(fold f init) (x:xs) = f x ((fold f init) xs)
```

where the parentheses around `fold f init` are used only for emphasis, and are otherwise superfluous.

From this we can now derive new definitions for `sum` and `product`:

```
sum = fold add 0
prod = fold mul 1
```

Of course now that the `fold` abstraction has been made, many other useful functions can be defined, even something as seemingly unrelated as `append`:

```
append xs ys = fold (:) ys xs
```

[An infix operator may be passed as an argument by surrounding it in parentheses; thus `(:)` is equivalent to `\x y -> x:y.`] This version of `append` simply replaces the `[]` at the end of the list `xs` with the list `ys`.

It is easy to verify that the new definitions are equivalent to the old using simple equational reasoning and induction. It is also important to note that in arriving at the main abstraction we did nothing out of the ordinary—we are just applying classical data abstraction principles in as unrestricted a way as possible, and that means allowing functions to be “first-class citizens.”

## 3.2 Non-Strict Semantics (Lazy Evaluation)

### 3.2.1 Fundamentals

The normal-order reduction rules of the lambda calculus are the most general in that they are guaranteed to produce a normal form if in fact one exists (see Section 2.1). In other words, they result in termination of the rewriting process most often, and the strategy lies at the heart of the Church-Rosser Theorem. Furthermore, as argued earlier, normal-order reduction allows recursion to be emulated with the Y combinator, thus giving the lambda calculus the most powerful form of effective computability, captured in Church's Thesis.

Given all this, it is quite natural to consider using normal-order reduction as the computational basis for a programming language. Unfortunately, normal-order reduction, implemented naively, is hopelessly inefficient. To see why, consider this simple normal-order reduction sequence:

$$\begin{aligned} & (\lambda x. (+ x x)) (* 5 4) \\ \Rightarrow & (+ (* 5 4) (* 5 4)) \\ \Rightarrow & (+ 20 (* 5 4)) \\ \Rightarrow & (+ 20 20) \\ \Rightarrow & 40 \end{aligned}$$

Note that the multiplication  $(* 5 4)$  is done *twice*. In the general case this could be an arbitrarily large computation, and it could potentially be repeated as many times as there are occurrences of the formal parameter (for this reason an analogy is often drawn between normal-order reduction and *call-by-name* parameter passing in Algol). In practice this can happen quite often, reflecting the simple fact that results are often shared.

One solution to this problem is to resort to something other than normal-order reduction, such as applicative-order reduction, which for the above term yields the following reduction sequence:

$$\begin{aligned} & (\lambda x. (+ x x)) (* 5 4) \\ \Rightarrow & (\lambda x. (+ x x)) 20 \\ \Rightarrow & (+ 20 20) \\ \Rightarrow & 40 \end{aligned}$$

Note that the argument is evaluated *before* the  $\beta$ -reduction is performed (similar to a *call-by-value* parameter-passing mechanism), and thus the normal form is reached in three steps instead of four, with no recomputation. The problem with this solution is that it requires the introduction of a special reduction rule to implement recursion (such as gained through the  $\delta$ -rule for McCarthy's conditional), and furthermore there

are examples where it does *more* work than normal-order reduction! For example, consider:

<p>applicative-order:</p> $\begin{aligned} & (\lambda x. 1) (* 5 4) \\ \Rightarrow & (\lambda x. 1) 20 \\ \Rightarrow & 1 \end{aligned}$	<p>normal-order:</p> $\begin{aligned} & (\lambda x. 1) (* 5 4) \\ \Rightarrow & 1 \end{aligned}$
--	--

or even worse (repeated from Section 2.1):

<p>applicative-order:</p> $\begin{aligned} & (\lambda x. 1) ((\lambda x. x x) (\lambda x. x x)) \\ \Rightarrow & (\lambda x. 1) ((\lambda x. x x) (\lambda x. x x)) \\ \Rightarrow & \vdots \end{aligned}$	<p>normal-order:</p> $\begin{aligned} & (\lambda x. 1) ((\lambda x. x x) (\lambda x. x x)) \\ \Rightarrow & 1 \end{aligned}$
--	--

which in the applicative-order case does not terminate! Despite these problems, most of the early functional languages, including pure Lisp, FP, ML, Hope, all of the dataflow languages, and others, used applicative-order semantics.<sup>19</sup> In addition to overcoming the efficiency problem of normal-order reduction, applicative-order reduction could be implemented with relative ease using the call-by-value compiler technology that had been developed for conventional imperative programming languages.

Nevertheless, the appeal of normal-order reduction cannot be ignored. Returning to lambda calculus basics, we can try to get to the root of the efficiency problem, which seems to be the following: reduction in the lambda calculus is normally described as *string* reduction, which precludes any possibility of sharing. If instead we were to describe it as a *graph* reduction process, perhaps sharing could be achieved. This idea was first suggested by Wadsworth in his PhD thesis in 1971 [Wad71, Chapter 4], in which he outlined a graph reduction strategy that utilized pointers to implement sharing. Using this strategy results in the following reduction sequence for the first example given earlier:

$$\begin{aligned} & (\lambda x. (+ x x)) (* 5 4) \\ \Rightarrow & (+ \bullet \bullet) \\ & \quad \swarrow \quad \searrow \\ & \quad \underbrace{\hspace{2em}} \\ & \quad (* 5 4) \\ \Rightarrow & (+ \bullet \bullet) \\ & \quad \swarrow \quad \searrow \\ & \quad \underbrace{\hspace{2em}} \\ & \quad 20 \\ \Rightarrow & 40 \end{aligned}$$

which takes the same number of steps as the applicative-order reduction sequence.

<sup>19</sup>Actually this is not quite true—most implementations of these languages use an applicative-order reduction strategy for the *top-level redices only*, thus yielding what is known as a “weak head normal form.” This strategy turns out to be easier to implement than complete applicative-order reduction, and also permits certain versions of the Y combinator to be implemented without special “rules.” See [Bur75] for an example of this using Landin’s SECD machine.

We will call an implementation of normal-order reduction in which recomputation is avoided *lazy evaluation* (another term often used is *call-by-need*). Its key feature is that arguments in function calls are evaluated *at most once*. It possesses the full power of normal-order reduction while being more efficient than applicative-order reduction in that “at most once” sometimes amounts to no computation at all!

Despite the appeal of lazy evaluation and this apparent solution to the efficiency problem, it took a good ten years more for researchers to discover ways to implement it efficiently compared to conventional programming languages. The chief problem has to do with the difficulty in implementing lazy graph reduction mechanisms on conventional computers, which seem to be better suited to call-by-value strategies. Simulating the unevaluated portions of the graph in a call-by-value environment amounts to efficiently implementing *closures*, or *thunks*, which have some inherent, non-trivial costs [BHY88]. It is beyond the scope of this paper to discuss the details of these implementation concerns, but see Peyton-Jones [PJ87] for an excellent summary.

Rather than live with conventional computers, one could alternatively build specialized graph reduction or dataflow hardware, but so far this has not resulted in any practical, much less commercially available, machines. Nevertheless, this work is quite promising, and good summaries of work in this area can be found in articles by Treleaven et al. [TBH82] and Vegdahl [Veg84], both of which are reprinted in [Tha87].

### 3.2.2 Expressiveness

Assuming that we can implement lazy evaluation efficiently (current practice is considered acceptably good), we should return to the question of why we want it in the first place. Previously we argued on philosophical grounds—it is the most general evaluation policy—but is lazy evaluation useful to programmers in practice? The answer is an emphatic yes, which I hope to show via a two-fold argument.

First of all, *lazy evaluation frees a programmer from concerns about evaluation order*. The fact is, programmers are generally concerned about the efficiency of their programs, and thus they prefer not evaluating things that are not absolutely necessary. As a simple example, suppose we *may* need to know the greatest common divisor of  $b$  and  $c$  in some computation involving  $x$ . In a modern functional language we might write:

```
f a x
where a = gcd b c
```

without worrying about  $a$  being evaluated needlessly—if in the computation of “ $f a x$ ” the value of  $a$  is needed, it will be computed, otherwise not. If we were to have written this program in Scheme, for example, we might try:



```
(let ( (a (gcd b c)) )
      (f a x))
```

which will *always* evaluate *a*. Knowing that *f* doesn't always need that value, and being concerned about efficiency, we may decide to rewrite this as:

```
(let ( (a (delay (gcd b c))) )
      (f a x))
```

which requires modifying *f* so as to force its first argument. Alternatively we could just write `(f b c x)`, which requires modifying *f* so as to compute the gcd internally. Both of these solutions are severe violations of modularity, and they arise out of the programmer's concern about evaluation order. Lazy evaluation eliminates that concern and preserves modularity.

The second argument for lazy evaluation is perhaps the one more often heard: *the ability to compute with unbounded ("infinite") data structures*. The idea of lazily evaluating data structures was first proposed by Vuillemin [Vui74], but similar ideas were developed independently by Henderson and Morris [HM76], and Friedman and Wise [FW76]. In a later series of papers [Tur81, Tur82] Turner provides a strong argument for using lazy lists, especially when combined with list comprehensions (see Section 2.7) and higher-order functions (see Section 3.1). Aside from Turner's elegant examples, Hughes presents an argument [Hug84] based on facilitating modularity in programming, where, along with higher-order functions, lazy evaluation is described as a way to "glue" pieces of programs together.

The primary power of lazily evaluated data structures comes from its utility in separating *data* from *control*. The idea is that a programmer should be able to describe a specific data structure without worrying about how it gets evaluated. Thus, for example, one could describe the sequence of natural numbers by the following simple program:

```
nats = 0 : map succ nats
```

or alternatively by:

```
numsfrom n = n : numsfrom (n+1)
nats       = numsfrom 0
```

These are examples of "infinite lists," or streams, and in a language that did not support lazy evaluation would cause the program to diverge. With lazy evaluation these data structures are only evaluated as they are needed, "on demand." For example, we could define a function that filters out only those elements satisfying a property *p*:

```
filter p (x:xs) = if (p x) then (x:rest) else rest
                  where rest = filter p xs
```

in which case “`filter p nats`” could be written knowing that the degree of the list’s computation will be determined by its context—i.e. the consumer of the result. Thus `filter` has no operational control within it, and can be combined with other functions in a modular way. For example one could compose it with a function to square each element in a list:

```
map (\x->x*x) . filter p
```

[ `.` is the infix composition operator.] This kind of modular result, in which data is separated from control, is one of the key features of lazy evaluation. Many more examples of this kind may be found in [Hug84].

### 3.3 Data Abstraction

Independently of the development of functional languages there has been considerable work on data abstraction in general, and on strong typing, user-defined datatypes, and type checking in particular. Some of this work has also taken on a theoretical flavor, not only in foundational mathematics where logicians have used types to resolve many famous paradoxes, but also in formal semantics where types aid our understanding of programming languages.

Fueling the theoretical work are two significant pragmatic advantages of using data abstraction and strong typing in one’s programming methodology. First, *data abstraction* improves the modularity, security, and clarity of programs. Modularity is improved because one can abstract away from implementation (i.e. representation) details; security is improved because interface violations are automatically prohibited; and clarity is improved because data abstraction has an almost self-documenting flavor.

Second, *strong static typing* helps in debugging since one is guaranteed that if a program compiles successfully that no error can occur at run-time due to type violations. It also leads to more efficient implementations, since it allows one to eliminate most run-time tag bits and type testing. Thus there is little performance penalty for using data abstraction techniques.

Of course, these issues are true for any programming language, and for that reason a thorough treatment of types and data abstraction is outside the scope of this survey; the reader may find an excellent summary in [CW85]. The basic idea behind the Hindley-Milner type system was discussed in Section 2.6.1. I will concentrate in this section on how *data abstraction* is manifest in modern functional languages.

### 3.3.1 Concrete Datatypes

As mentioned earlier, there is a strong argument for wanting language features that facilitate data abstraction, whether or not the language is functional. In fact such mechanisms were first developed in the context of imperative languages such as Simula, Clu, Euclid, and others. It is only natural that they be included in functional languages. ML, as we mentioned, was the first functional language to do this, but many others soon followed suit.

In this subsection I will describe *concrete* (or *algebraic*) datatypes as well as *type synonyms*. I will use Haskell syntax, but the ideas are essentially identical (at least semantically) to those used in ML and Miranda.

New algebraic datatypes may be defined along with their constructors using data declarations, as in the following definitions of lists and trees:

```
data List a = Nil   | Cons a (List a)
data Tree b = Empty | Node b (List (Tree b))
```

The identifiers `List` and `Tree` are called *type constructors*, and the identifiers `a` and `b` are called *type variables* which are implicitly universally quantified over the scope of the data declaration. The identifiers `Nil`, `Cons`, `Empty`, and `Node` are called *data constructors*, or just *constructors*, with `Nil` and `Empty` being *nullary constructors*. [Note that both type constructors and data constructors are capitalized, so that the latter can be used without confusion in pattern-matching (as discussed in the next section) and to avoid confusion with type variables (such as `a` and `b` in the above example).]

`List` and `Tree` are called type constructors since they construct types from other types. For example, `Tree Ints` is the type of trees of integers. Reading from the data declaration for `Tree`, we see then that a tree of integers is either `Empty`, or a `Node` containing a list of more trees of integers.

We can now see that the previously given type signature for `map`:

```
map :: (a -> b) -> [a] -> [b]
```

is equivalent to:

```
map :: (a -> b) -> (List a) -> (List b)
```

That is, “[...]” in a type expression is just syntax for application of the type constructor `List`. Similarly, we can think of `->` as an infix type constructor that creates the type of all functions from its first argument (a type) to its second (also a type).

[A useful property to note is the consistent syntax used in Haskell for expressions and types. Specifically, if  $T_i$  is the type of expression or pattern  $e_i$ , then the expressions  $\backslash e_1 \rightarrow e_2$ ,  $[e_1]$ , and  $(e_1, e_2)$  have the types  $T_1 \rightarrow T_2$ ,  $[T_1]$ , and  $(T_1, T_2)$ , respectively.]

Instances of these new types are built simply by using the constructors. Thus “Empty” is an empty tree, and “Node 5 Empty” is a very simple tree of integers with one element. The type of the instance is inferred via the same type inference mechanism that infers types of polymorphic functions, as described previously.

Defining new concrete datatypes is fairly common not only in functional languages but also in imperative languages, although the polymorphism offered by modern functional languages makes it all the more attractive.

*Type synonyms* are a way of creating new names for types, such as in the following:

```
type Intree    = Tree Ints
type Flattener = Intree -> [Ints]
```

Note that Intree is used in the definition of Flattener. `type` declarations do not introduce *new* types (as `data` declarations do), but rather are a convenient way to introduce new *names* (i.e. synonyms) for existing types.

### 3.3.2 Abstract Datatypes

Another idea in data abstraction originating in imperative languages is the notion of an *abstract datatype* (ADT) in which the details of the implementation of a datatype are hidden from the users of that type, thus enhancing modularity and security. The traditional way to do this is exemplified by ML’s ADT facility and emulated in Miranda. Although the Haskell designers chose a different approach to ADT’s (described below), the following example of a queue ADT is written as if Haskell had ML’s kind of ADT’s, using the keyword `abstype`:

```
abstype Queue a = Q [a]
where first (Q as) = last as
      isempty (Q []) = True
      isempty (Q as) = False
      ...
```

The main point is that the functions `first`, `isempty`, etc. are visible in the scope of the `abstype` declaration, *but the constructor* `Q`, *including its type, is not*. Thus a user of the ADT has no idea whether queues are implemented as lists (as shown here) or some other data structure. The advantage of this, of course, is that one is free to change the representation type without fear of “breaking” some other code that uses the ADT.

### 3.3.3 Haskell's Orthogonal Design

In Haskell a rather different approach was taken to ADT's. The observation was made that the main difference between a concrete and abstract datatype was that the latter had a certain degree of *information hiding* built into it. So instead of thinking of abstract datatypes and information hiding as going hand-in-hand, the two were made orthogonal components of the language. More specifically, concrete datatypes were made the only data abstraction mechanism, and to that an `expose` declaration was added to control information hiding, or "visibility."

For example, to get the effect of the earlier definition of a queue, one would write:

```
expose Queue, first, isempty
from data Queue a = Q [a]
    first (Q as) = last as
    isempty (Q []) = True
    isempty (Q as) = False
    ...
```

Since `Q` is not explicitly listed in the `expose` declaration, it becomes hidden from the user of the ADT.

The advantage of this approach to ADT's is more flexibility. For example, suppose we also wish to hide `isempty`, or perhaps some auxiliary function defined in the nested scope. This is trivially done with the orthogonal design but much harder with the ML design as described so far. Indeed, to alleviate this problem the ML designers provided an additional construct, a `local` declaration, with which one can hide local declarations. Another advantage of the orthogonal design is that the same mechanism can be used at the top level of a module to control visibility of the internals of the module to the external world. In other words, the `expose` mechanism is very general, and can be nested. Haskell utilizes a conservative module system that relies on this capability.

A disadvantage of the orthogonal approach is that if the most typical ADT scenario only requires hiding the representation type, the user will have to think through the details in each case rather than having the hiding done automatically by using `abstype`.

## 3.4 Equations and Pattern-Matching

One of the programming methodology attributes that is strongly encouraged in the modern school of functional programming is the use of *equational reasoning* in the design and construction of programs. The lack of side effects accounts for the primary ability to apply equational reasoning, but there are syntactic features that can facilitate it as well. Using *equations* as part of the syntax is the most obvious of these, but along

with this goes *pattern-matching*, in which one can write several equations in defining the same function, only one of which is presumably applicable in a given situation. Thus modern functional languages have tried to maximize the expressiveness of pattern-matching.

At first blush, equations and pattern-matching seem fairly intuitive and relatively innocuous. Indeed, we have already given many examples that use pattern-matching, without having said much about the details of how it works. But in fact pattern-matching can have surprisingly subtle effects on the semantics of a language, and thus should be carefully and precisely defined.

### 3.4.1 Pattern-Matching Basics

Pattern-matching should actually be viewed as the primitive behavior of a *case expression*, which has the general form:

```
case e of
  pat1 -> e1
  pat2 -> e2
  ...
  patn -> en
```

Intuitively, if the structure of  $e$  matches  $pat_i$  then the result of the case expression is  $e_i$ . A set of equations of the form:

```
f pat1 = e1
f pat2 = e2
...
f patn = en
```

can then be thought of as shorthand for:

```
f = \x -> case x of
  pat1 -> e1
  pat2 -> e2
  ...
  patn -> en
```

Despite this translation, for convenience I will use the equational syntax in the remainder of this section.

The question to be asked first is just what the pattern-matching process consists of. For example, what exactly are we pattern-matching *against*? One idea that seems reasonable is to allow one to pattern-match against constants and data structures. Thus `fac` can be defined by:

```

fac 0 = 1
'   n = n*fac(n-1)

```

[The “tick mark” in the second equation is an abbreviation for `fac`.] But note that this relies on a “top-to-bottom” reading of the program, since `(fac 0)` actually matches both equations. We can remove this ambiguity by adding a *guard* (recall the discussion in Section 2.7), which in Haskell looks like:

```

fac 0      = 1
'   n | n>0 = n*fac(n-1)

```

As we have already demonstrated through several examples, it is also reasonable to pattern-match against lists:

```

length [] = 0
'   (x:xs) = 1 + length xs

```

and for that matter any data structure, including user-defined ones:

```

data Tree2 a = Leaf a | Branch (Tree2 a) (Tree2 a)
fringe (Leaf x)          = [x]
'   (Branch left right) = fringe left ++ fringe right

```

where `++` is the infix append operator.

Another idea that seems desirable is the ability to repeat formal parameters on the left-hand-side to denote that arguments in those positions must have the same value, as in the second line of the following definition:

```

member x [] = False
'   x (x:xs) = True
'   x (y:xs) = member x xs

```

[This is not legal Haskell syntax, since such repetition of identifiers is not allowed.] However, care must be taken with this approach since in something like:

```

alleq [x,x,x] = True
'   y          = False

```

it is not clear in what order the elements of the list are to be evaluated. For example, if they are evaluated left-to-right then `alleq [1,2,bot]`, where `bot` is any non-terminating computation, will return `False`, whereas with a right-to-left order the program will diverge. One solution that would at least guarantee consistency in this approach is to insist that all three positions are evaluated, so that the program diverges if any of the arguments diverge.

In general the problem of what gets evaluated, and when, is perhaps the most subtle aspect of reasoning about pattern-matching, and suggests that the pattern-matching algorithm be fairly simple so as not to mislead the user. Thus in Haskell the above repetition of identifiers is disallowed—equations must be *linear*—but some functional languages allow it (for example Miranda and Alfl [Hud84]).

A particularly subtle version of this problem is captured in the following example. Consider these definitions:

```
data Silly a = Foo a | Other
bar (Foo x) = 0
'   Other   = 1
```

Then a call `bar bot` will diverge, since `bar` must be strict in order that it can distinguish between the two kinds of arguments that it might receive. But now consider this small modification:

```
data Silly a = Foo a
bar (Foo x) = 0
```

Then a call `bar bot` seems like it should return 0, since `bar` need not be strict—it can only receive one kind of argument, and thus does not need to examine it unless it is needed in computing the result, which in this case it is not. Given the goal of a language that is “lazy as possible,” this seems like the right solution, and was the approach adopted in Haskell.

Two useful discussions on the subject of pattern-matching can be found in [Aug85] and [Wad].

### 3.4.2 Connecting Equations

Let us now turn to the more global issue of how the individual equations are connected together as a group. As mentioned earlier, one simple way to do this is give the equations a top-to-bottom priority, so that in:

```
fac 0 = 1
'   n = n*(fac(n-1))
```

the second equation is tried only after the first one has failed. This is the solution adopted in many functional languages, including Haskell and Miranda.

An alternative method is to insist that the equations be disjoint, thus rendering the order irrelevant. One significant motivation for this is the desire to reason about



the applicability of an equation independently of the others, thus facilitating equational reasoning. The question is, how can one guarantee disjointness? For equations without guards, the disjointness property can be determined statically; i.e. by just examining the patterns. Unfortunately, when unrestricted guards are allowed, the problem obviously becomes undecidable, since it amounts to determining the equality of arbitrary recursive predicates. This in turn can be solved by resolving the guard disjointness at run-time. On the other hand, this solution ensures correctness only for values actually encountered at run-time, and thus the programmer might apply equational reasoning erroneously to as yet un-encountered values!

The two ideas could also be combined by providing two different syntaxes for joining equations. For example, using the hypothetical keyword `else` (not valid in Haskell):

```

sameShallowStructure [a]   [c]   = True
',
                    [a,b] [c,d] = True
else
',
                    x     y     = False

```

The first two equations would be combined using a disjoint semantics; together they would then be combined with the third using a top-to-bottom semantics. Thus the third equation acts as an “otherwise” clause in the case that the first two fail. A design of this sort was considered for Haskell early on, but the complexities of disjointness, especially in the context of guards, were considered too great and the design was eventually rejected.

### 3.4.3 Argument Order

In the same sense that it is desirable to have the order of equations be irrelevant, it is also desirable to have the order of *arguments* be irrelevant. In exploring this possibility, consider first the functions `f` and `g` defined by:

```

f 1 1 = 1
f 2 x = 2

g 1 1 = 1
g x 2 = 2

```

which differ only in the order of their arguments. Now consider what “`f 2 bot`” should evaluate to. Clearly the equations for `f` are disjoint, clearly the expression matches only the second equation, and since we want a non-strict language it seems the answer should clearly be 2. For a compiler to achieve this it must always evaluate the *first* argument to `f` first.

Now consider the expression “`g bot 2`”—by the same argument given above the result should also be 2, but now the compiler must be sure to always evaluate the *second* argument to `g` first. Can a compiler always determine the correct order in which to evaluate its arguments?

To help answer this question, first consider this intriguing example (due to Berry [Ber78]):

$$\begin{aligned}f\ 0\ 1\ x &= 1 \\f\ 1\ x\ 0 &= 2 \\f\ x\ 0\ 1 &= 3\end{aligned}$$

Clearly these equations are disjoint. So what is the value of “`f 0 1 bot`”? Desired answer: 1. And what is the value of “`f 1 bot 0`”? Desired answer: 2. And what is the value of “`f bot 0 1`”? Desired answer: 3. But now the most difficult question: In what order should the arguments be evaluated? If we evaluate the third one first, then the answer to the first question cannot be 1. If we evaluate the second one first, then the answer to the second question cannot be 2. If we evaluate the first one first, then the answer to the third question cannot be 3. In fact there is *no* sequential order that will allow us to answer these three questions the way we would like—some kind of parallel evaluation is required.

This subtle problem is solvable in several ways, but they all require some kind of compromise—insisting on a parallel (or pseudo-parallel) implementation, rejecting certain seemingly valid programs, making equations more strict than one would like, or giving up on the independence of the order of evaluation of arguments. In Haskell the last solution was chosen—perform a left-to-right evaluation of the arguments—because it presented the simplest semantics to the user, which was judged to be more important than making the order of arguments irrelevant. Another way to explain this is to think of equations as syntax for nested lambda expressions, in which case one might not expect symmetry with respect to the arguments anyway.

### 3.5 Formal Semantics

Simultaneously with work on functional languages Scott, Strachey, and others were busy establishing the foundations of denotational semantics, now the most widely used tool for describing the formal semantics of programming languages. There was a close connection between this work and functional languages primarily because the lambda calculus served as one of the simplest programming languages with enough useful properties to make its formal semantics interesting. In particular, the lambda calculus had a notion of *self-application*, which implied that certain domains had to contain their own function spaces. That is, it required a domain  $D$  that was a solution

to the following domain equation:

$$D = D \rightarrow D$$

At first this seems impossible—surely there are more functions from  $D$  into  $D$  than there are elements in  $D$ —but Scott was able to show that indeed such domains existed, as long as one was willing to restrict the allowable functions in certain (quite reasonable) ways, and by treating  $=$  as an *isomorphism* rather than an equality [Sco70]. Scott’s work served as the mathematical foundation for Strachey’s work [MS76] on the denotational semantics of programming languages; see [Sto77] and [Sch85] for thorough treatments.

Denotational semantics and functional programming have close connections, and the functional programming community emphasizes the importance of formal semantics in general. For completeness, and to show how simple the denotational semantics of a functional language can be, we give the semantics of the *recursive lambda calculus with constants* defined in Section 2.3.

$$\begin{aligned} Bas &= Int + Bool + \dots && \text{Basic values} \\ D &= Bas + (D \rightarrow D) && \text{Denotable values} \\ Env &= Id \rightarrow D && \text{Environments} \end{aligned}$$

$$\begin{aligned} \mathcal{E} &: Exp \rightarrow Env \rightarrow D \\ \mathcal{K} &: Con \rightarrow D \end{aligned}$$

$$\begin{aligned} \mathcal{E}[[x]]env &= env[[x]] \\ \mathcal{E}[[c]]env &= \mathcal{K}[[c]] \\ \mathcal{E}[[e_1 e_2]]env &= (\mathcal{E}[[e_1]]env) (\mathcal{E}[[e_2]]env) \\ \mathcal{E}[[\lambda x.e]]env &= \lambda v. \mathcal{E}[[e]]env[v/x] \\ \mathcal{E}[[e \text{ where } x_1 = e_1; \dots; x_n = e_n]]env &= \mathcal{E}[[e]]env' \\ &\text{where } env' = \text{fix } \lambda env'. env[ (\mathcal{E}[[e_1]]env') / x_1, \\ &\quad \vdots \\ &\quad (\mathcal{E}[[e_n]]env') / x_n ] \end{aligned}$$

This semantics is relatively simple, but in moving to a more complex language such as Miranda or Haskell the semantics can become significantly more complex, due to the many syntactic features that make the languages convenient to use. In addition, one must state precisely the *static* semantics as well, including type checking, pattern-matching usage, etc. Once all is said and done, the formal semantics of Haskell consumes about 20 pages.

## 4 Advanced Features and Active Research Areas

Some of the most recent ideas in functional language design are new enough that they should be regarded as on-going research. Nevertheless many of them are sound enough to have been included in current language designs. In this section we will explore a variety of such ideas, beginning with some of the innovative ideas in the Haskell design. Some of the topics have a theoretical flavor, such as extensions of the Hindley-Milner type system; some have a pragmatic flavor, such as expressing non-determinism, efficient arrays, and I/O; and some involve the testing of new application areas, such as parallel and distributed computation. All in all studying these topics should provide insight into the goals of functional programming as well as some of the problems in achieving those goals.

### 4.1 Overloading

The kind of polymorphism exhibited by the Hindley-Milner type system is what Strachey called *parametric* polymorphism, to distinguish it from another kind that he called *ad hoc* polymorphism, or *overloading*. The two can be distinguished in the following way: a function with parametric polymorphism *does not care* what type certain of its arguments have, and thus it behaves identically regardless of the type. In contrast, a function with *ad hoc* polymorphism *does* care, and in fact may behave differently for different types. Stated another way, *ad hoc* polymorphism is really just a syntactic device for overloading a particular function name or symbol with more than one meaning.

For example, the function *map* defined earlier exhibits parametric polymorphism, and has typing:

$$\text{map} :: (a \rightarrow b) \rightarrow [a] \rightarrow [b]$$

Regardless of the kind of list given to *map* it behaves in the same way. In contrast, consider the function *+* which we normally wish to behave differently for integer and floating point numbers, and not at all (i.e. be a static error) for non-numeric arguments. Another common example is the function *==* (equality) which certainly behaves differently when comparing the equality of two numbers versus, say, two lists.

*Ad hoc* polymorphism is normally (and I suppose appropriately!) treated in an *ad hoc* manner. Worse, there is no accepted convention for doing this; indeed ML, Miranda, and Hope all do it differently. Recently, however, a uniform treatment for handling overloading was discovered independently by Kaes [Kae88] (in trying to generalize ML's *ad hoc* "equality types") and Wadler and Blott (as part of the process of defining Haskell). Below I will describe the solution as adopted in Haskell; details may be found in [WB89].

The basic idea is to introduce a notion of *type classes* that capture a collection of overloaded operators in a consistent way. A class declaration is used to introduce a new type class and the overloaded operators that must be supported by any type that is an instance of that class. An instance declaration declares that a certain type is an instance of a certain class, and thus included in the declaration are the definitions of the overloaded operators instantiated on the named type.

For example, say that we wish to overload `+` and `negate` on types `Int` and `Float`. To do so, we introduce a new type class called `Num`:

```
class Num a where
  (+)      :: a -> a -> a
  negate  :: a -> a
```

This declaration may be read “a type `a` belongs to the class `Num` if there are (overloaded) functions `+` and `negate`, of the appropriate types, defined on it.”

We may then declare `Int` and `Float` to be instances of this class, as follows:

```
instance Num Int where
  x + y      = addInt x y
  negate x   = negateInt x

instance Num Float where
  x + y      = addFloat x y
  negate x   = negateFloat x
```

[Note how infix operators are defined; Haskell’s lexical syntax prevents ambiguities.] where `addInt`, `negateInt`, `addFloat`, and `negateFloat` are assumed in this case to be pre-defined functions, but in general could be any user-defined function. The first declaration above may be read “`Int` is an instance of the class `Num` as witnessed by these definitions of `+` and `negate`.”

Using type classes one can thus treat overloading in a consistent, arguably elegant, way. Another nice feature is that type classes naturally support a notion of *inheritance*. For example, we may define a class `Eq` by:

```
class Eq a where
  (==) :: a -> a -> Bool
```

Given this class, we would certainly expect all members of the class `Num`, say, to have `==` defined on them. Thus the class declaration for `Num` could be changed to:

```
class Eq a => Num a where
  (+)      :: a -> a -> a
  negate  :: a -> a
```

which can be read as “only members of the class `Eq` may be members of the class `Num`, and a type `a` belongs to the class `Num` if ... <as before>.” Given this class declaration, instance declarations for `Num` must include a definition of `==`, as in:

```
instance Num Int where
  x + y      = addInt x y
  negate x   = negateInt x
  x == y     = eqInt x y
```

The Haskell Report uses this inheritance mechanism to define a very rich hierarchical numeric structure which reflects fairly well a mathematician’s view of numbers.

The traditional Hindley-Milner type system is extended in Haskell to include type classes. The resulting type system is able to verify that the overloaded operators do have the appropriate type. However, it is possible (but not likely) for ambiguous situations to arise, which in Haskell result in type error but can be reconciled explicitly by the user (see [HWe88] for the details).

## 4.2 Purely Functional Yet Universal I/O

To many the notion of I/O conjures an image of state, side-effects, and sequencing. Is there any hope at achieving *purely functional*, yet *universal*, and of course *efficient* I/O? Surprisingly, the answer is yes. Perhaps even more surprising is that over the years there have emerged not one, but *two* seemingly very different solutions:

- The *lazy stream* model, in which temporal events are modelled as lists, whose lazy semantics mimics the demand-driven behavior of processes.
- The *continuation* model in which temporality is modelled via explicit continuations.

Although papers have been written advocating both solutions, and indeed they are very different in *style*, the two solutions turn out to be exactly equivalent in terms of *expressiveness*; in fact there is an almost trivial translation from one to the other. The Haskell I/O system takes advantage of this fact and provides a unified framework that supports both styles. The specific I/O operations available in each style are identical—what differs is the way they are expressed—and thus programs in either style may be combined with a well-defined semantics. In addition, although certain of the primitives rely on non-deterministic behavior in the operating system, referential transparency is still retained *internal* to a Haskell program.

In this section the two styles will be described as they appear in Haskell, together with the translation of one in terms of the other. Details of the actual Haskell design



```
type Behavior = [Response] -> [Request]
```

[Recall from Section 3.3 that [Response] is the type consisting of lists of values of type Response.] The main operational idea is that the *n*th response is the reply of the operating system to the *n*th request.

For simplicity we will assume that there are only two kinds of requests and three kinds of responses, as defined below:

```
data Request = ReadFile Name | WriteFile Name Contents
data Response = Success | Return Contents | Failure ErrorMsg
type Name = String
type Contents = String
type ErrorMsg = String
```

[This is a subset of the requests available in Haskell.]

As an example, given this request list:

```
[ ..., WriteFile fname s1, Readfile fname, ... ]
```

and the corresponding response list:

```
[ ..., Success, Return s2, ... ]
```

then  $s1 == s2$ , unless there were some “intervening” external effect.

In contrast, the continuation model is normally characterized by a set of *transactions*. Each transaction typically takes a *success continuation* and a *failure continuation* as arguments. These continuations in turn are simply functions that generate more transactions. For example:

```
data Transaction = ReadFile' Name          FailCont RetCont
                 | WriteFile' Name Contents FailCont SuccCont
                 | Done
type FailCont = ErrorMsg -> Transaction
type RetCont  = Contents -> Transaction
type SuccCont =          Transaction
```

[In Haskell the transactions are actually provided as *functions* rather than *constructors*; see below.] The special transaction Done represents program termination. These declarations should be compared with those for the stream model given earlier.

Returning to the simple example given earlier, the request and response list are no longer separate entities since their effect is “interwoven” into the continuation structure, yielding something like:



```

WriteFile' fname s1 exit
  (ReadFile' fname exit
   (\s2 -> ...) )
where exit errmsg = Done

```

in which case, as before, we would expect `s1 == s2` in the absence of external effects. This is essentially the way I/O is handled in Hope.

Although these two styles seem very different, there is a very simple translation of the continuation model into the stream model. In Haskell, instead of defining the new datatype `Transaction`, a set of *functions* are defined that accomplish the same task, but that are really stream transformers in the request/response style. In other words, the type `Transaction` should be precisely the type `Behavior`, and should not be a new datatype at all. Thus we arrive at:

```

readFile  :: Name ->          FailCont -> RetCont  -> Behavior
writeFile :: Name -> Contents -> FailCont -> SuccCont -> Behavior
done      ::                               Behavior

type FailCont = ErrorMessage -> Behavior
type RetCont  = Contents     -> Behavior
type SuccCont =              Behavior

readFile name fail succ resps =
  (ReadFile name) : case (head resps) of
                    Return contents -> succ contents (tail resps)
                    Failure msg    -> fail msg (tail resps)
writeFile name contents fail succ resps =
  (WriteFile name contents) : case (head resps) of
                               Success      -> succ (tail resps)
                               Failure msg -> fail msg (tail resps)

done resps = []

```

This pleasing and very efficient translation allows one to write Haskell programs in either style and mix them freely.

The complete design, of course, includes a fairly rich set of primitive requests besides `ReadFile` and `WriteFile`, including a set of requests for communicating through *channels*, which include things such as standard input and output. One of these requests, in fact, takes a *list* of input channel names and returns their *non-deterministic merge*. Because the merge itself is implemented in the operating system, referential transparency is retained within a Haskell program.

Another useful aspect of the Haskell design is that it includes a partial (but rigorous) specification of the behavior of the OS itself. For example, it introduces the notion of

an *agent* that consumes data on output channels and produces data on input channels. The *user* is then modelled as an agent that consumes standard output and produces standard input. This particular agent is required to be *strict* in the standard output, corresponding to the notion that the user reads the terminal display before typing at the keyboard! No other language design that I am aware of has gone this far in specifying the I/O system with this degree of precision; it is usually left implicit in the specification. It is particularly important, however, in this context because the proper semantics relies critically on how the OS consumes and produces the request and response lists.

To conclude this section I will show two Haskell programs that perform the following task: prompt the user for the name of a file, read and echo the file name, and then look up and display the contents of the file on standard-output. The first version uses the stream model, the second the continuation model.

```
main resps =
  [ AppendChannel "stdout" "please type a filename\CR\",
    if (resps!!1 == Success) then (ReadChannel "stdin"),
    AppendChannel "stdout" fname,
    if (resps!!3 == Success) then (ReadFile fname),
    AppendChannel "stdout" (case resps!!4 of
                            Failure msg      -> "can't open"
                            Return file_contents -> file_contents)
  ] where fname = case resps!!2 of
                  Return user_input -> get_line user_input
```

[The operator !! is the list selection operator; thus xs!!n is the nth element in the list xs.]

```
main = appendChannel "stdout" "please type a filename\CR\" exit
      (readChannel "stdin" exit (\user_input ->
        appendChannel "stdout" fname exit
        (readFile fname (\msg -> appendChannel "stdout" "can't open" exit done)
          (\contents ->
            appendChannel "stdout" contents exit done))
        where fname = get_line user_input))
      exit msg = done
```

### 4.3 Arrays

As it turns out, arrays can be expressed rather nicely in a functional language, and in fact all of the arguments about mathematical elegance fall in line when using arrays.

This is especially true of program development in scientific computation, where textbook matrix algebra can often be translated almost verbatim into a functional language. The main philosophy is to treat the entire array as a single entity defined declaratively, rather than as a place-holder of values that is updated incrementally. This in fact is the basis of the APL philosophy (see Section 2.4), and some researchers have concentrated on combining functional programming ideas with those from APL [Tu86, TP86]. The reader may find good general discussions of arrays in functional programming languages in [Wad86, Hud86a, Wis87]. In the remainder of this section I will describe Haskell's arrays, which originated from some ideas in Id Nouveau [NPA86].

Haskell has a family of multi-dimensional non-strict immutable arrays whose special interaction with list comprehensions provides a convenient "array comprehension" syntax for defining arrays monolithically. As an example, here is how to define a vector of squares of the integers from 1 to  $n$ :

```
a = array (1,n) [ (i,i*i) | i <- [1..n] ]
```

The first argument to `array` is a tuple of *bounds*, and thus this array has size  $n$  and is indexed from 1 to  $n$ . The second argument is a list of index/value pairs, and is written here as a conventional list comprehension. The  $i$ th element of an array `a` is written `a!i`, and thus in the above case we have that `a!i = i*i`.

There are several useful semantic properties of Haskell's arrays. First, they can be recursive—here is an example of defining the first  $n$  numbers in the Fibonacci sequence:

```
fib = array (1,n) ( [ (0,1), (1,1) ] ++
  [ (i, fib!(i-1)+fib!(i-2)) | i <- [2..n] ] )
```

This example demonstrates how one can use an array as a *cache*, which in this case turns an exponentially poor algorithm into an efficient linear one.

Another important property is that array comprehensions are constructed lazily, and thus the order of the elements in the list is completely irrelevant. For example, we can construct a  $m$ -by- $n$  matrix using a "wavefront recurrence" where the north and west borders are 1, and each other element is the sum of its north, north-west, and west neighbors, as follows:

```
a = array ((1,m),(1,n))
  ( [ ((1,1),1) ] ++
    [ ((i,1),1) | i <- [2..m] ] ++
    [ ((1,j),1) | j <- [2..n] ] ++
    [ ((i,j), a!(i-1,j) + a!(i,j-1) + a!(i-1,j-1))
      | i <- [2..m], j <- [2..n] ] )
```

The elements in this result can be accessed in any order—the “demand-driven” effect of lazy evaluation will cause the necessary elements to be evaluated in an order constrained only by data dependencies. It is this property that makes array comprehensions so useful in scientific computation, where recurrence equations express the same kind of data dependencies. In implementing such recurrences in Fortran one must be sure that the elements are evaluated in an order consistent with the dependencies—lazy evaluation accomplishes that for us.

On the other hand, although elegant, array comprehensions may be difficult to implement efficiently. There are two main difficulties: the standard problem of overcoming the inefficiencies of lazy evaluation, and the problem of avoiding the construction of the many intermediate lists that the second argument to array seems to need. A discussion of ways to overcome these problems is found in [AH89]. Alternative designs for functional arrays and their implementations may be found in [Hol83, Hug85a, AHN87, Wis87].

Another problem is that array comprehensions are not quite expressiveness enough to capture all behaviors. The most conspicuous example of this is the case where an array is being used as an “accumulator,” say in building a histogram, and thus one actually *wants* an incremental update effect. Thus in Haskell a function called `accumArray` is provided to capture this kind of behavior in a way consistent with the “monolithic” nature of array comprehensions (similar ideas are found in [SH86, Wad86]). However, it is not clear that this is the most general solution to the problem. An alternative approach is to define an incremental update operator on arrays, but then even nastier efficiency problems arise, since (conceptually at least) the updates involve copying. Work on detecting when it is safe to implement such updates destructively has resulted in at least one efficient implementation [BH88, Blo88, HB85], although the analysis itself is costly.

Nevertheless, array comprehensions have the potential for being very useful, and many interesting applications have already been programmed using them (see [HA88] for some examples). Hopefully future research will lead to solutions to the remaining problems.

## 4.4 Views

Pattern-matching (see Section 3.4) is very useful in writing compact and readable programs. Unfortunately, knowledge of the concrete representation of an object is necessary before pattern-matching can be invoked, which seems to be at odds with the notion of an abstract datatype. To reconcile this conflict Wadler introduced a notion of *views* [Wad87].

A `view` declaration introduces a new algebraic datatype, just like a `data` declaration, but in addition establishes an isomorphism between the values of this new type and a subset of the values of an existing algebraic datatype. For example:

```

data Complex = Rectangular Float Float
view Complex = Polar      Float Float
  where toView  (Rectangular x y) = Polar (sqrt (x**2+y**2))
                                             (arctan (y/x))
        fromView (Polar r t) = Rectangular (r*(cos t))
                                             (r*(sin t))

```

[Views are not part of Haskell, but as with abstract datatypes we will use Haskell syntax, here extended with the keyword `view`.] Given the datatype `Complex`, we can read the view declaration as: “one view of `Complex` contains only `Polar` values; to translate from a `Rectangular` to a `Polar` value, use the function `toView`; to translate the other way, use `fromView`.”

Having declared this view we can now use pattern-matching using either the `Rectangular` constructor or the `Polar` constructor; similarly, new objects of type `Complex` can be built with either the `Rectangular` or `Polar` constructors. They have precisely the same status, and the coercions are done automatically. For example:

```
rotate (Polar r t) angle = Polar r (t+angle)
```

As the example stands, objects of type `Complex` are concretely represented with the `Rectangular` constructor, but this decision could be reversed by making `Polar` the concrete constructor and `Rectangular` the view, *without altering any of the functions which manipulate objects of type `Complex`*.

Whereas traditionally abstract data types are regarded as hiding the representation, with views one can reveal as many representations (zero, one, or more) as are required.

As a final example, consider this definition of “Peano’s view” of the natural number subset of integers:

```

view Integer = Zero | Succ Integer
  where fromView Zero      = 0
        ' (Succ n) | n>=0 = n+1
        toView 0          = Zero
        ' n | n>0 = Succ (n-1)

```

With this view, 7 is viewed as equivalent to

```
Succ (Succ (Succ (Succ (Succ (Succ (Succ Zero))))))
```

Note that `fromView` defines a mapping of any finite element of peano into an integer, and `toView` defines the inverse mapping. Given this view, we can write definitions such as

```
fac Zero    = 1
' (Succ n) = (Succ n) * (fac n)
```

which is very useful from an “equational reasoning” standpoint, since it allows us to use an abstract representation of integers *without* incurring any performance overhead—the view declarations provide enough information to map all of the abstractions back into concrete implementations at compile time.

On the other hand, perhaps the most displeasing aspect of views is that an implicit coercion is taking place, which may be confusing to the user. For example, in:

```
case (Foo a b) of
  Foo x y -> exp
```

one cannot be sure in `exp` that `a==x` and `b==y`! Although views were considered in an initial version of Haskell, they were eventually discarded, in a large part because of this problem.

## 4.5 Parallel Functional Programming

An often-heralded advantage of functional languages is that parallelism in a functional program is *implicit*; it is manifested solely through data dependencies and the semantics of primitive operators. This is in contrast to more conventional languages, where explicit constructs are typically used to invoke, synchronize, and in general coordinate the concurrent activities. In fact, as discussed earlier, many functional languages were developed simultaneously with work on highly-parallel dataflow and reduction machines, and such research continues today.

In most of this work parallelism in a functional program is detected by the system and allocated to processors *automatically*. Although in certain constrained classes of functional languages the mapping of process to processor can be determined *optimally* [Che86, DI85], in the general case the optimal strategy is undecidable, so heuristics such as load-balancing are often employed instead.

But what if a programmer knows a good (perhaps optimal) mapping strategy for a program executing on a particular machine, but the compiler is not smart enough to determine it? And even if the compiler is smart enough, how does one *reason* about such behavior? One could argue that the programmer should not be concerned about such details, but that is a difficult argument to make to someone whose job is precisely to invent such algorithms!

To meet these needs various researchers have designed extensions to functional languages, resulting in what I like to call *para-functional programming languages*. The extensions amount to a *meta-language* (for example annotations) to express the desired

behavior. Examples include annotations to control evaluation order [Bur84, DW87, Sri85], prioritize tasks, and map processes to processors [Hud86c, HS86, KL85]. Similar work has taken place in the Prolog community [Sha84]. In addition, research has resulted in formal operational semantics for such extensions [Hud86b, HA87]. In the remainder of this section one kind of para-functional behavior will be demonstrated: that of mapping program to machine (based on the work in [Hud86c, HS86]).

The fundamental idea behind process-to-processor mapping is quite simple. Consider the expression  $e1+e2$ . The strict semantics of  $+$  allows the subexpressions  $e1$  and  $e2$  to be executed in parallel—this is an example of what is meant by saying that “the parallelism in a functional program is implicit.” But suppose now that we wish to express precisely *where* (i.e., on which processor) the subexpressions are to be evaluated; we may do so quite simply by annotating the subexpressions with appropriate mapping information. An expression annotated in this way is called a *mapped expression*, which has the following form:

`exp on proc`

[`on` is a hypothetical keyword, and is not valid Haskell.] which intuitively declares that `exp` is to be computed on the processor identified by `proc`. The expression `exp` is the *body* of the mapped expression, and represents the value to which the overall expression will evaluate (and thus can be any valid expression, including another mapped expression). The expression `proc` must evaluate to a processor id. Without loss of generality the processor ids, or *pids*, are assumed to be *integers*, and there is some pre-defined mapping from those integers to the physical processors they denote.

Returning now to the example, we may annotate the expression  $(e1+e2)$  as follows:

`(e1 on 0) + (e2 on 1)`

where 0 and 1 are processor ids. Of course, this static mapping is not very interesting. It would be nice, for example, if we were able to refer to a processor *relative to the currently executing one*. We can do this through the use of the reserved identifier `self`, which when evaluated returns the pid of the currently executing processor. Using `self` we can now be more creative. For example, suppose we have a ring of  $n$  processors that are numbered consecutively; we may then rewrite the above expression as:

`(e1 on left self) + (e2 on right self)`  
 where `left pid = mod (pid-1) n`  
       `right pid = mod (pid+1) n`

[`mod x y` computes  $x$  modulo  $y$ .] which denotes the computation of the two subexpressions in parallel on the two neighboring processors, with the sum being computed on `self`.

To see that it is desirable to dynamically bind `self`, consider that one may wish successive invocations of a recursive call to be executed on different processors—this is not easily expressed with lexically bound annotations. For example, consider the following list-of-factorials program, again using a ring of processors:

```
(map fac [2,3,4]) on 0
where map f [] = []
      ' f (x:xs) = f x : ((map f xs) on (right self))
```

Note that the recursive call to `map` is mapped onto the processor to the right of the current one, and thus the elements 2, 6, and 24 in the result list are computed on processors 0, 1, and 2, respectively.

Para-functional programming languages have been shown to be adequate in expressing a wide range of deterministic parallel algorithms clearly and concisely [Hud86c, HS86]. It remains to be seen, however, whether the pragmatic concerns that motivate these kinds of language extensions persist, and if they do, whether or not compilers can become smart enough to perform the optimizations automatically. Indeed these same questions can be asked about other language extensions, such as the “memoization” techniques discussed in the next section.

## 4.6 Caching and Memoization

Consider this simple definition of the Fibonacci function:

```
fib 0 = 1
' 1 = 1
' n = fib (n-1) + fib (n-2)
```

Although simple, it is hopelessly inefficient. One could rewrite it in one of the classic ways, but then the simplicity and elegance of the original definition is lost. Keller and Sleep suggest an elegant alternative: provide syntax for expressing the *caching* or *memoization* of selected functions [KS86]. For example, the syntax might take the form of a declaration that precedes the function definition, as in:

```
memo fib using cache
fib 0 = 1
' 1 = 1
' n = fib (n-1) + fib (n-2)
```

which would be syntactic sugar for:



```

fib = cache fib1
  where fib1 0 = 1
         '    1 = 1
         '    n = fib (n-1) + fib (n-2)

```

The point is that `cache` is a user-defined function that specifies a strategy for caching values of `fib`. For example, to cache values in an array, one might define `cache` by:

```

cache fn = \n -> (array (0,max) [(i,fn i) | i<-[0..max]]) ! n

```

where we assume `max` is the largest argument to which `fib` will be applied. Expanding out the definitions and syntax yields:

```

fib n = (array (0,max) [(i,fib1 i) | i<-[0..max]]) ! n
  where fib1 0 = 1
         '    1 = 1
         '    n = fib (n-1) + fib (n-2)

```

which is exactly the desired result.<sup>20</sup> As a methodology this is very nice, since libraries of useful caching functionals may be produced and reused to cache many different functions. There are limitations to the approach, as well as extensions, all of which are described in [KS86].

One of the limitations of this approach is that in general it can only be used with strict functions, and even then the expense of performing equality checks on, for example, list arguments can be expensive. As a solution to this problem, Hughes introduced the notion of “lazy memo-functions,” in which the caching strategy utilizes an *identity* test (“EQ” in Lisp terminology) instead of an equality test [Hug85b]. Such a strategy can no longer be considered as syntactic sugar, since an identity predicate is not something normally provided as a primitive in functional languages because it is implementation dependent. Nevertheless, if built into a language lazy memo-functions provide a very efficient (constant time) caching mechanism, and allow very elegant solutions to a class of problems not solved by Keller and Sleep’s strategy: those involving “infinite” data structures. For example, consider this definition of the infinite list of ones:

```

ones = 1 : ones

```

Any reasonable implementation will represent this as a cyclic list, thus consuming constant space. Another common idiom is the use of higher order functions such as `map`:

```

twos = map (\x->2*x) ones

```

---

<sup>20</sup>This is essentially the same solution as the one given in Section 4.3.

But now note that only the cleverest of implementations will represent this as a cyclic list, since `map` normally generates a new list cell on every recursive call. By lazy memoizing `map`, however, the cyclic list will be recovered. To see how, note that the first recursive call to `map` will be `map (\x->2*x) (tail ones)` — but `(tail ones)` is *identical* to `ones` (remember that `ones` is cyclic), and thus `map` is called with arguments identical to the first call. Thus the old value is returned, and the cycle is created. Many more practical examples are discussed in [Hug85b].

The interesting thing about memoization in general is that it begins to touch on some of the limitations of functional languages—in particular the inability to side effect global objects such as caches—and solutions such as lazy memo-functions represent useful compromises. It remains to be seen whether more general solutions can be found that eliminate the need for these special-purpose features.

## 4.7 Non-Determinism

Most programmers (including the very idealistic among us) admit the need for non-determinism, despite the semantic difficulties that it introduces. It seems to be an essential ingredient of “real-time systems,” such as operating systems and device controllers. Non-determinism in imperative languages is typically manifested by running in parallel several processes that are side-effecting some global state—the non-determinism is thus *implicit* in the semantics of the language. In functional languages non-determinism is manifested through the use of primitive operators such as *amb* or *merge*—the non-determinism is thus made *explicit*. Several papers have been published on the use of such primitives in functional programming, and it appears quite reasonable to program conventional non-deterministic applications using them [Hen82, Sto85]. The problem is, once introduced, non-determinism completely destroys referential transparency, as we shall see.

By way of introduction, McCarthy [McC63] defined a binary non-deterministic operator called *amb* having the following behavior:

$$\begin{aligned} \text{amb}(e_1, \perp) &= e_1 \\ \text{amb}(\perp, e_2) &= e_2 \\ \text{amb}(e_1, e_2) &= \text{either } e_1 \text{ or } e_2, \text{ chosen non-deterministically} \end{aligned}$$

The “operational reading” of  $\text{amb}(e_1, e_2)$  is that  $e_1$  and  $e_2$  are evaluated in parallel, and the one that completes first is returned as the value of the expression.

To see how referential transparency is lost, consider this simple example:

$$(\text{amb } 1 \ 2) + (\text{amb } 1 \ 2)$$

Is the answer 2 or 4? Or perhaps 3? It is this last answer that indicates that referential transparency is lost—there does not appear to be any simple syntactic mechanism for

ensuring that one could not replace “equals-for-equals” in a misleading way. Shortly we will discuss possible solutions to this problem, but for now let us look at an example using this style of non-determinism.

Using `amb` one can easily define things such as `merge` that non-deterministically merge two lists, or *streams*:<sup>21</sup>

```
merge as bs = amb (if (as==[]) then bs else (head as : merge (tail as) bs))
                (if (bs==[]) then as else (head bs : merge as (tail bs)))
```

which then can be used, for example, in combining streams of characters from different computer terminals:

```
process (merge term1 term2)
```

Using this as a basis, Henderson shows how many operating system problems can be solved in a “pseudo-functional” language [Hen82]. Hudak uses non-determinism of a slightly different kind to emulate the parallel updating of arrays [Hud86a].

Although satisfying in the sense of being able to solve real-world kinds of non-determinism, these solutions are dissatisfying in the way they destroy referential transparency. One might argue that the situation is at least somewhat better than the conventional imperative one in that the non-determinism is at least made *explicit*, and thus one could induce extra caution when reasoning about those sections of a program exhibiting non-deterministic behavior. The only problem with this is that determining which sections of a program are non-deterministic may be difficult—it is not a *lexical* property, but rather is *dynamic*, since any function may call a non-deterministic sub-function.

At least two solutions have been proposed to this problem. One, proposed by Warren Burton [Bur88], is to provide a tree-shaped “oracle” as an argument to a program from which non-deterministic choices may be selected. By passing the tree and its sub-trees around explicitly, referential transparency can be preserved. The problem with this approach is that carrying the oracle around explicitly is cumbersome at best. On the other hand, functional programmers already carry around a greater amount of state to avoid problems with side-effects, so perhaps the extra burden is not too great!

Another (at least partial) solution was proposed by Stoye [Sto85] in which all of the non-determinism in a program is forced to be in one place. Indeed to some extent this

---

<sup>21</sup>Note that this version of `merge`:

```
merge [] bs = bs
merge as [] = as
merge (a:as) (b:bs) = amb (a : merge as (b:bs)) (b : merge (a:as) bs)
```

is not correct, since “`merge [] bs`” evaluates to “`⊥`,” whereas we would like it to be “`bs ++ ⊥`,” which in fact the definition in the text yields.

was the solution adopted in Haskell, although for somewhat different reasons. The problem with this approach is that the non-determinism is not eliminated completely, but rather centralized. It allows reasoning equationally within the isolated pieces, but not within the collection of pieces as a whole. Nevertheless, the isolation (at least in Haskell) is achieved *syntactically*, and thus it is easy to determine when equational reasoning is valid. A general discussion of these issues is found in [HS88].

An interesting variation on these two ideas is to *combine* them—centralize the non-determinism and then use an oracle to define it in a referentially transparent way. Thus the disadvantages of both approaches would seem to disappear.

In any case, I should point out that *none* of these solutions makes reasoning about non-determinism any easier, they just make reasoning about *programs* easier.

#### 4.8 Extensions to Polymorphic Type Inference

The Hindley-Milner type system has certain limitations; an example of this was given in Section 2.6.1. Some success has been achieved in extending the type system to include other kinds of data objects, but surprisingly little success has been achieved at removing the fundamental limitations while still retaining the tractability of the type inference problem. It is a somewhat fragile system, but is fortunately expressive enough to base practical languages on it. Nevertheless, research continues in this area.

Independently of type inference, considerable research is underway on the expressiveness of type systems in general. The most obvious thing to do is allow types to be “first-class,” thus allowing abstraction over them in the obvious way. Through generalizations of the type system it is possible to model such things as parameterized modules, inheritance, and sub-typing. This area has indeed taken on a character of its own; a good summary of current work may be found in both [CW85] and [Rey85].

#### 4.9 Combining Other Programming Language Paradigms

A time-honored tradition in programming language design is to come up with “hybrid” designs that combine the best features of several different paradigms, and functional programming language research has not escaped that tradition! I will only discuss briefly two such hybrids here, although others exist.

The first is combining logic programming with functional programming. The “logical variable” permits two-way matching (via unification) in logic programming languages, as opposed to the one-way matching (via pattern-matching) in functional languages, and thus seems like a desirable feature to have in a language. Indeed its declarative nature fits in well with the ideals of functional programming. However the integration is not easy—many proposals have been made yet none are completely

satisfactory, especially in the context of higher-order functions and lazy evaluation. See [DL85] for a good summary of results.

The second area reflects an attempt to combine the state-oriented behavior of imperative languages in a semantically clean way. The same problems arise here as they do with non-determinism—simply adding the assignment statement means that equational reasoning must always be qualified, since it is difficult to determine whether or not a deeply nested call is made to a function that induces a side-effect. However, there are some possible solutions to this, most notably work by Gifford and Lucassen [GL86, LG88] in which effects are captured in the *type system*. In Gifford’s system it is possible to determine from a procedure’s type whether or not it is side-effect free. However, it is not currently known whether such a system can be made into a type inference system in which type declarations are not required. This is an active area of current research.

## 5 Dispelling Myths About Functional Programming

To gain further insight into the nature of functional languages, it is helpful to discuss, with the hope of dispelling, certain myths that have arisen over the years.

**Myth #1:** Functional programming is the antithesis of conventional imperative programming.

This myth is largely responsible for alienating imperative programmers from functional languages. But in fact, there is much in common between the two styles of programming, which I hope to make evident by two simple arguments.

Consider first that one of the key evolutionary characteristics of high-level imperative programming languages has been the use of *expressions* to denote a result rather than a sequence of statements to put together a result imperatively and in piecemeal. Expressions were an important feature of Fortran, for example, had more of a mathematical flavor, and freed the programmer of low-level operational detail (this burden was of course transferred to the compiler). From Fortran expressions, to functions in Pascal, to expression oriented programming style in Scheme—these advances are all on the same evolutionary path. Functional programming can be seen as carrying this evolution to its logical conclusion—*everything* is an expression.

The second argument is based on an analogy between functional (i.e. side-effect-free) programming and structured (i.e. goto-less) programming. The fact is, it is hard to imagine doing without either goto’s or assignment statements, *until one is shown what to use in their place*. In the case of goto, one uses instead “structured commands,” and in the case of assignment statements one uses instead lexical binding and recursion.

As an example, this simple program fragment with goto’s:

```

x := init;
i := 0;
loop: x := f(x,i);
      i := i+1;
      if i<10 goto loop;

```

can be rewritten in a structured style as:

```

x := init;
i := 0;
while i<10
  begin x := f(x,i);
        i := i+1
  end;

```

In capturing this disciplined use of goto, arbitrary jumps into or out of the body of the block now cannot be made. Although this can be viewed as a constraint, most people now feel that the resulting *disciplined* style of programming is clearer, easier to maintain, etc.

But in fact more discipline is evident here than just the judicious use of goto. Note in the original program fragment that *x* and *i* are assigned to exactly *once* in each iteration of the loop; the variable *i*, in fact, is only being used to control the loop termination criteria, and the final value of *x* is intended as the value computed by the loop. This *disciplined use of assignment* can be captured by the following assignment-free Haskell program:

```

loop init 0
where loop x i = if i<10 then (loop (f x i) (i+1))
                  else x

```

Functions (and procedures) can in fact be thought of as a disciplined use of goto *and* assignment—the transfer of control to the body of the function and the subsequent return capture a disciplined use of goto, and the formal to actual parameter binding captures a disciplined use of assignment.

By inventing a bit of syntactic sugar to capture the essence of tail recursion, the above program could be rewritten as:

```

let x = init
    i = 0
in while i<10
  begin next x = f(x,i)

```

```
        next i = i+1
    end
    result x
```

[This syntactic sugar is not found in Haskell, although some other functional (especially dataflow) languages have similar features, including `Id`, `Val`, and `Lucid`.] where the form “`next x = ...`” is a construct (`next` is a keyword) used to express what the value of `x` will be on the *next* iteration of the “loop.” Note the similarity of this program to the structured one given earlier. In order to properly enforce a “disciplined” use of assignment, we can constrain the syntax so that only one “`next`” statement is allowed for each identifier (stated another way, this constraint means that it is a trivial matter to convert such programs into the tail recursive form shown earlier). If one thinks of “`next x`” and “`next i`” as *new* identifiers (just as the formal parameters of `loop` can be thought of as new identifiers for every call to `loop`), then referential transparency is preserved. Functional programming advocates argue that this results in a better style of programming, in much the same way that structured programming advocates argue for their cause.

Thus the analogy between goto-less programming and assignment-free programming runs deep. When Dijkstra first introduced structured programming, much of the programming community was aghast—how could one do without `goto`? But as people programmed in the new style, it was realized that what was being imposed was a “discipline” for good programming, not a “police state” to inhibit expressiveness. Exactly the same can be said of side-effect-free programming, and its advocates hope that as people become more comfortable programming in the functional style, they will appreciate the good sides of the discipline thus imposed.

When viewed in this way functional languages can be seen as a logical step in the evolution of imperative languages—thus, of course, rendering them non-imperative. On the other hand, it is exactly this purity that some programmers object to, and one could argue that just as a “tasteful” use of `goto` here or there is acceptable, so is a “tasteful” use of a side-effect. Such small impurities certainly shouldn’t invalidate the functional programming *style*, and thus may be acceptable.

**Myth #2:** Functional programming languages are toys.

The first step toward dispelling this myth is to cite examples of efficient implementations of functional languages, of which there now exists several. The `Alfl` compiler at Yale, for example, generates code that is competitive with that generated by conventional language compilers [You88]. Other notable compiler efforts include the `LML` compiler at Chalmers University [Aug84], the `Ponder` compiler at Cambridge University [Fai85], and the `ML` compilers developed at Bell Labs and Princeton [AM87].

On the other hand, there are still inherent inefficiencies that cannot be ignored. Higher-order functions and lazy evaluation certainly increase expressiveness, but in the

general case the overhead of, for example, the dynamic storage management necessary to support them, cannot be eliminated.

The second step toward dispelling this myth amounts to pointing to real applications development in functional languages, including “real-world” situations involving non-determinism, databases, parallelism, etc. Although examples of this sort are not plentiful (primarily because of the youth of the field) and are hard to cite (since papers are not usually written about applications), they do exist. For example, I know of several:

1. The dataflow groups at MIT and the functional programming groups at Yale have written numerous functional programs for scientific computation. So has two national labs: Los Alamos and Lawrence Livermore.
2. MCC has written a reasonably large expert system (EMYCIN) in SASL.
3. At least one company is currently marketing a commercial product (a CAD package) that uses a lazy functional language.
4. A group at IBM uses a lazy functional language for graphics and animation.
5. The LML (lazy ML) compiler at Chalmers was written almost entirely in LML, and the new Haskell compilers at both Glasgow and Yale are being written in Haskell.
6. GEC Hirst Research Lab is using a program for designing VLSI circuits that was written by some researchers at Oxford using a lazy functional language.

I am sure that there are many other examples that I am not aware of. In particular, there are many Scheme and Lisp programs that are predominantly side-effect-free, and could properly be thought of as functional programs.

**Myth #3:** Functional languages cannot deal with state.

This criticism is often expressed as a question: “how can one program in a language that does not have a notion of state?” The answer, of course, is that you cannot, and in fact functional languages deal with state very nicely, although the state is expressed *explicitly* rather than *implicitly*. So the issue is more a matter of how one *expresses* state and manipulations of it.

State in a functional program is usually carried around explicitly in one of two ways: (1) in the values of bound variables of functions, in which case it is updated by making a recursive call to the function with new values as arguments, or (2) in a data structure that is updated non-destructively so that, at least conceptually, the “old” value of the data structure remains intact and can be accessed later. Although declarative and referentially transparent, this treatment of state can present problems to an implementation, but it is certainly not a problem in expressiveness. Furthermore, the



implementation problems are not as bad as they first seem, and recent work has gone a long way toward solving them. For example, a compiler can convert tail recursions into loops and “single-threaded” data structures into mutable ones.

It turns out that, with respect to expressiveness, one can use higher-order functions not only to manipulate state, but also to make it appear implicit. To see how, consider the imperative program given earlier:

```
x := init;
i := 0;
loop: x := f(x,i);
      i := i + 1;
      if (i < 10) goto loop;
```

We can model the implicit state in this program as a pair  $(xval, ival)$ , and define several functions that update this state in a way that mimics the assignment statements:

```
x (xval,ival) xval' = (xval',ival)
i (xval,ival) ival' = (xval,ival')
x' (x,i) = x
i' (x,i) = i
const v s = v
```

We will take advantage of the fact that these functions are defined in curried form.

Note how  $x$  and  $i$  are used to *update* the state, and  $x'$  and  $i'$  are used to *access* the state. For example, the following function, when applied to the state, will increment  $i$ :

```
\s -> i s ((i' s) + 1)
```

For expository purposes we would like to make the state as implicit as possible, and thus we express the result as a *composition* of higher-order functions. To facilitate this, and to make the result look as much like the original program as possible, we define the following higher-order infix operators and functions:<sup>22</sup>

```
f := g = \s -> f s (g s)
f ; g = \s -> g (f s)
goto f = f
f +' g = \s -> f s + g s
f <' g = \s -> f s < g s
if' p c = \s -> (if (p s) then (c s) else s)
```

<sup>22</sup>It is interesting to note that `:=`, `const`, and `goto` correspond precisely to the combinators  $S$ ,  $K$ , and  $I$ , and `;` is almost the combinator  $B$ , but is actually  $CB$ .

[I am cheating slightly here in that `;` is a reserved operator in Haskell, and thus cannot really be redefined in this way.]

Given these definitions we can now write the following functional (albeit contrived) version of the imperative program given earlier:

```
      x := const init;
      i := const 0;
loop where
loop = x := f;
      i := i' + const 1;
      if' (i' < const 10) (goto loop)
```

This result is rather disquieting—it looks very much like the original imperative program! Of course, we worked hard to make it that way, which in the general case is much harder to do, and it is certainly not the recommended way to do functional programming. Nevertheless it exemplifies the power and flexibility of higher-order functions—note how they are used here both to manipulate the state and to implement the `goto` (where in particular the definition of `loop` is recursive, since the `goto` implements a loop).

## 6 Conclusions

In this paper I have attempted to present functional programming in all of its many shapes and forms. In so doing I have only touched on the surface of many issues. It is hoped that enough of a foundation has been given that the interested researcher may explore particularly interesting topics in more depth, and that the interested programmer may learn to use functional languages in a variety of applications.

The reader will note that I have said very little about how to *implement* functional languages, primarily because doing that subject justice would probably double the size of this paper! Nevertheless, I feel compelled to point the interested reader in the right direction for such a study. By far the best single reference to *sequential* implementations is Peyton Jones' text [PJ87], although recently other techniques have appeared viable (see [BHY88, BPJR88, FW87]). Parallel implementations have taken a variety of forms. On commercial machines the state of the art on parallel graph reduction implementations may be found in [GH88, Gol88b, Gol88a]. The latest on special-purpose parallel graph reducers can be found in [PJCSH87, WW87]. For a very different kind of implementation see [HM88]. References to dataflow machines were given in Section 2.8.

**Acknowledgements.** Thanks to the Lisp and Functional Programming Research Group at Yale, who inspire much of my work and serve as my chief sounding board. A

special thanks is also extended to the Haskell Committee, through which I learned more than I care to admit. In addition, the following people provided valuable comments on earlier drafts of the paper: Kei Davis, Alex Ferguson, John Launchbury, and Phil Wadler (all from the University of Glasgow); Juan Guzman, Siau-Cheng Khoo, Amir Kishon, Raman Sundaresh, and Pradeep Varma (all from Yale); David Wise (Indiana University); and three anonymous referees,

I also wish to thank my funding agencies: The Department of Energy under grant FG02-86ER25012, the Defense Advanced Research Projects Agency under grant N00014-88-K-0573, and the National Science Foundation under grant DCR-8451415. Without their generous support none of this work would have been possible.

In writing this survey paper I have tried to acknowledge accurately the significant technical contributions in each of the areas of functional programming that I have covered. Unfortunately, that is a very difficult task, and I apologize in advance for any errors or omissions.

## References

- [AD79] W.B. Ackerman and J.B. Dennis. *VAL - a value-oriented algorithmic language preliminary reference manual*. Laboratory for Computer Science MIT/LCS/TR-218, MIT, June 1979.
- [AG77] Arvind and K.P. Gostelow. A computer capable of exchanging processors for time. In *Proceedings IFIP Congress*, pages 849–853, June 1977.
- [AG82] Arvind and K.P. Gostelow. The U-interpreter. *Computer*, 15(2):42–50, February 1982.
- [AH89] S. Anderson and P. Hudak. *Efficient Compilation of Haskell Array Comprehensions*. Technical Report YALEU/DCS/RR693, Yale University, Department of Computer Science, March 1989.
- [AHN87] A. Aasa, S. Holmstrom, and C. Nilsson. *An Efficiency Comparison of Some Representations of Purely Functional Arrays*. Technical Report 33, Programming Methodology Group, Chalmers University of Technology, May 1987.
- [AK81] Arvind and V. Kathail. A multiple processor data flow machine that supports generalized procedures. In *Proceedings 8th Annual Symposium Computer Architecture*, pages 291–302, ACM SIGARCH 9(3), May 1981.
- [AM87] A.W. Appel and D.B. MacQueen. A standard ML compiler. In *Proceedings of 1987 Functional Programming Languages and Computer Architecture Conference*, pages 301–324, Springer-Verlag LNCS 274, September 1987.
- [ASS85] H. Abelson, G.J. Sussman, and J. Sussman. *Structure and Interpretation of Computer Programs*. The MIT Press (Cambridge, MA) and McGraw-Hill Book Company (New York, NY), 1985.
- [Aug84] L. Augustsson. A compiler for Lazy ML. In *Proceedings 1984 ACM Conference on LISP and Functional Programming*, pages 218–227, August 1984.
- [Aug85] L. Augustsson. Compiling pattern-matching. In *Functional Programming Languages and Computer Architecture*, pages 368–381, Springer-Verlag LNCS 201, September 1985.
- [AW76a] E.A. Ashcroft and W.W. Wadge. Lucid - a formal system for writing and proving programs. *SIAM Journal of Computing*, 5(3):336–354, September 1976.
- [AW76b] E.A. Ashcroft and W.W. Wadge. Lucid, a nonprocedural language with iteration. *Communications of the ACM*, 20:519–526, 1976.

- [Bac78] J. Backus. Can programming be liberated from the von Neumann style? A functional style and its algebra of programs. *CACM*, 21(8):613-641, August 1978.
- [Bar84] H.P. Barendregt. *The Lambda Calculus, Its Syntax and Semantics*. North-Holland, Amsterdam, 1984. Revised edition.
- [BD77] R.M. Burstall and J. Darlington. A transformation system for developing recursive programs. *JACM*, 24(1):44-67, 1977.
- [Ber78] G. Berry. Séquentialité de l'évaluation formelle des  $\lambda$ -expressions. In *Proceedings 3-e Colloque International sur la Programmation*, March 1978.
- [BH88] A. Bloss and P. Hudak. Path semantics. In *Proceedings of Third Workshop on the Mathematical Foundations of Programming Language Semantics*, Springer-Verlag LNCS (to appear), 1988.
- [BHY88] A. Bloss, P. Hudak, and J. Young. Code optimizations for lazy evaluation. *Lisp and Symbolic Computation: An International Journal*, 1:147-164, 1988.
- [Blo88] A. Bloss. *Path Analysis: Using Order-of-Evaluation Information to Optimize Lazy Functional Languages*. PhD thesis, Yale University, Department of Computer Science, 1988.
- [BMS80] R.M. Burstall, D.B. MacQueen, and D.T. Sannella. HOPE: an experimental applicative language. In *The 1980 LISP Conference*, pages 136-143, Stanford University, August 1980.
- [Boe85] H-J. Boehm. Partial polymorphic type inference is undecidable. In *Proceedings of 26th Symposium on Foundations of Computer Science*, pages 339-345, October 1985.
- [Bou88] B.E. Boutel. *Tui Language Manual*. Technical Report CSD-8-021, Victoria University of Wellington, Department of Computer Science, 1988.
- [BPJR88] G.L. Burn, S.L. Peyton Jones, and J.D. Robson. The spineless G-machine. In *Proceedings 1988 ACM Conference on Lisp and Functional Programming*, pages 244-258, ACM SIGPLAN/SIGACT/SIGART, Salt Lake City, Utah, August 1988.
- [Bur75] W.H. Burge. *Recursive Programming Techniques*. Addison-Wesley, Reading, MA, 1975.
- [Bur84] F.W. Burton. Annotations to control parallelism and reduction order in the distributed evaluation of functional programs. *ACM Transactions on Programming Languages and Systems*, 6(2):159-174, April 1984.

- [Bur88] F.W. Burton. Nondeterminism with referential transparency in functional programming languages. *The Computer Journal*, 31(3):243-247, 1988.
- [BW88] R. Bird and P. Wadler. *Introduction to Functional Programming*. Prentice Hall, New York, 1988.
- [BWW86] J. Backus, J.H. Williams, and E.L. Wimmers. *FL Language Manual (Preliminary Version)*. Technical Report RJ 5339 (54809) Computer Science, IBM Almaden Research Center, November 1986.
- [Car76] R. Cartwright. *A Practical Formal Semantic Definition and Verification System for Typed Lisp*. Technical Report, Stanford Artificial Intelligence Laboratory AIM-296, 1976.
- [CF58] H.B. Curry and R. Feys. *Combinatory Logic, Vol. 1*. North-Holland, Amsterdam, 1958.
- [Che86] M.C. Chen. Transformations of parallel programs in crystal. In *Information Processing '86*, pages 455-462, IFIP, Elsevier Science Publishers B.V. (North-Holland), 1986.
- [Chu33] A. Church. A set of postulates for the foundation of logic. *Annals of Mathematics*, 2(33-34):346-366,839-864, 1932-1933.
- [Chu41] A. Church. *The Calculi of Lambda Conversion*. Princeton University Press, Princeton, NJ, 1941.
- [CR36] A. Church and J.B. Rosser. Some properties of conversion. *Transactions of the American Mathematical Society*, 39:472-482, 1936.
- [Cur30] H.B. Curry. Grundlagen der kombinatorischen logik. *American Journal of Mathematics*, 52:509-536,789-834, 1930.
- [CW85] L. Cardelli and P. Wegner. On understanding types, data abstraction, and polymorphism. *Computing Surveys*, 17(4):471-522, December 1985.
- [Dav78] A.L. Davis. The architecture and system method of DDM-1: a recursively-structured data driven machine. In *Proceedings Fifth Annual Symposium on Computer Architecture*, 1978.
- [DI85] J.-M. Delosme and I.C.F. Ipsen. An illustration of a methodology for the construction of efficient systolic architectures in VLSI. In *Proceedings 2nd International Symposium on VLSI Technology, Systems, and Applications*, pages 268-273, 1985.
- [DL85] D. Degroot and G. Lindstrom. *Functional and Logic Programming*. Prentice-Hall, 1985.

- [DM74] J.B. Dennis and D.P. Misunas. A preliminary architecture for a basic data-flow processor. In *Proceedings of the 2nd Annual Symposium on Computer Architecture*, pages 126–132, ACM, IEEE, 1974.
- [DM82] L. Damas and R. Milner. Principle type schemes for functional languages. In *9th ACM Symposium on Principles of Programming Languages*, ACM, August 1982.
- [DW87] J. Darlington and L. While. Controlling the behavior of functional language systems. In *Proceedings of 1987 Functional Programming Languages and Computer Architecture Conference*, pages 278–300, Springer-Verlag LNCS 274, September 1987.
- [Fai85] J. Fairbairn. *Design and Implementation of a Simple Typed Language Based on the Lambda Calculus*. PhD thesis, University of Cambridge, May 1985. Available as Computer Laboratory TR No. 75.
- [FH88] A.J. Field and P.G. Harrison. *Functional Programming*. Adison-Wesley, Worthingam, England, 1988.
- [FLO85] S. Fortune, D. Leivant, and M. O’Donnell. The expressiveness of simple and second-order type structures. *JACM*, 30(1):151–185, January 1985.
- [FW76] D.P. Friedman and D.S. Wise. Cons should not evaluate its arguments. In *Automata, Languages and Programming*, pages 257–284, Edinburgh University Press, 1976.
- [FW87] J. Fairbairn and S. Wray. Tim: a simple, lazy abstract machine to execute supercombinators. In *Proceedings of 1987 Functional Programming Languages and Computer Architecture Conference*, pages 34–45, Springer Verlag LNCS 274, September 1987.
- [GH88] B. Goldberg and P. Hudak. Implementing functional programs on a hypercube multiprocessor. In *Proceedings of Third Conference on Hypercube Concurrent Computers and Applications*, ACM, January 1988.
- [GHG60] H. Gelernter, J.R. Hansen, and C.L. Gerberich. A FORTRAN-compiled list processing language. *JACM*, 7(2):87–101, 1960.
- [GHW81] J. Guttag, J. Horning, and J. Williams. FP with data abstraction and strong typing. In *Proceedings of the 1981 Conference on Functional Programming Languages and Computer Architecture*, pages 11–24, ACM, 1981.
- [Gir72] J.-Y. Girard. *Interprétation Fonctionnelle et Elimination des Coupures dans l’Arithmétique d’Ordre Supérieur*. PhD thesis, Univ. of Paris, 1972.

- [GL86] D.K. Gifford and J.M. Lucassen. Integrating functional and imperative programming. In *Proceedings 1986 ACM Conference on Lisp and Functional Programming*, pages 28–38, ACM SIGPLAN/SIGACT/SIGART, August 1986.
- [GMM\*78] M. Gordon, R. Milner, L. Morris, M. Newey, and C. Wadsworth. A metalanguage for interactive proof in LCF. In *Conference Record of the Fifth Annual ACM Symposium on Principles of Programming Languages*, pages 119–130, ACM, 1978.
- [GMW79] M.J. Gordon, R. Milner, and C.P. Wadsworth. *Edinburgh LCF*. Springer-Verlag LNCS 78, Berlin, 1979.
- [Gol88a] B. Goldberg. Buckwheat: graph reduction on a shared memory multiprocessor. In *Proceedings 1988 ACM Conference on Lisp and Functional Programming*, pages 40–51, ACM SIGPLAN/SIGACT/SIGART, Salt Lake City, Utah, August 1988.
- [Gol88b] B. Goldberg. *Multiprocessor Execution of Functional Programs*. PhD thesis, Yale University, Department of Computer Science, 1988. Available as technical report YALEU/DCS/RR-618.
- [HA87] P. Hudak and S. Anderson. Pomset interpretations of parallel functional programs. In *Proceedings of 1987 Functional Programming Languages and Computer Architecture Conference*, pages 234–256, Springer Verlag LNCS 274, September 1987.
- [HA88] P. Hudak and S. Anderson. *Haskell Solutions to the Language Session Problems at the 1988 Salishan High-Speed Computing Conference*. Technical Report YALEU/DCS/RR-627, Yale University, Department of Computer Science, January 1988.
- [Han] P. Hancock. Polymorphic type-checking. Chapters 8 and 9 in [PJ87].
- [HB85] P. Hudak and A. Bloss. The aggregate update problem in functional programming systems. In *12th ACM Symposium on Principles of Programming Languages*, pages 300–314, ACM, 1985.
- [Hen80] P. Henderson. *Functional Programming: Application and Implementation*. Prentice-Hall, Englewood Cliffs, NJ, 1980.
- [Hen82] P. Henderson. Purely functional operating systems. In *Functional Programming and Its Applications: An Advanced Course*, pages 177–192, Cambridge University Press, 1982.
- [Hin69] R. Hindley. The principle type scheme of an object in combinatory logic. *Transactions of the American Mathematical Society*, 146:29–60, December 1969.



- [HM76] Henderson and Morris. A lazy evaluator. In *3rd ACM Symposium on Principles of Programming Languages*, pages 95–103, ACM, January 1976.
- [HM88] P. Hudak and E. Mohr. Graphinators and the duality of SIMD and MIMD. In *Proceedings 1988 ACM Conference on Lisp and Functional Programming*, pages 224–234, ACM SIGPLAN/SIGACT/SIGART, Salt Lake City, Utah, August 1988.
- [Hol83] S. Holmstrom. How to handle large data structures in functional languages. In *Proc. of SERC/Chalmers Workshop on Declarative Programming Languages*, 1983.
- [HS86] P. Hudak and L. Smith. Para-functional programming: a paradigm for programming multiprocessor systems. In *12th ACM Symposium on Principles of Programming Languages*, pages 243–254, January 1986.
- [HS88] P. Hudak and R. Sundares. *On the Expressiveness of Purely Functional I/O Systems*. Technical Report YALEU/DCS/RR-665, Yale University, Department of Computer Science, December 1988.
- [Hud84] P. Hudak. *ALFL Reference Manual and Programmer's Guide*. Research Report YALEU/DCS/RR-322, Second Edition, Yale University, October 1984.
- [Hud86a] P. Hudak. Arrays, non-determinism, side-effects, and parallelism: a functional perspective. In *Proceedings of the Santa Fe Graph Reduction Workshop*, pages 312–327, Los Alamos/MCC, Springer-Verlag LNCS 279, October 1986.
- [Hud86b] P. Hudak. Denotational semantics of a para-functional programming language. *International Journal of Parallel Programming*, 15(2):103–125, April 1986.
- [Hud86c] P. Hudak. Para-functional programming. *Computer*, 19(8):60–71, August 1986.
- [Hug84] J. Hughes. *Why Functional Programming Matters*. Technical Report 16, Programming Methodology Group, Chalmers University of Technology, November 1984.
- [Hug85a] J. Hughes. *An Efficient Implementation of Purely Functional Arrays*. Technical Report, Programming Methodology Group, Chalmers University of Technology, 1985.
- [Hug85b] J. Hughes. Lazy memo-functions. In *Functional Programming Languages and Computer Architecture*, pages 129–146, Springer-Verlag LNCS 201, September 1985.

- [HWe88] P. Hudak and P. Wadler (editors). *Report on the Functional Programming Language Haskell*. Technical Report YALEU/DCS/RR656, Yale University, Department of Computer Science, November 1988.
- [Ive62] K. Iverson. *A Programming Language*. Wiley, New York, 1962.
- [Joh88] S.D. Johnson. *Daisy Programming Manual*. Technical Report, Indiana University Computer Science Department, 1988. (in progress, draft available on request).
- [Kae88] S. Kaes. Parametric polymorphsim. In *Proceedings of the 2nd European Symposium on Programming*, Springer-Verlag LNCS 300, March 1988.
- [Kel82] R.M. Keller. *FEL programmer's guide*. AMPS TR 7, University of Utah, March 1982.
- [KJRL80] R.M. Keller, B. Jayaraman, D. Rose, and G. Lindstrom. *FGL programmer's guide*. AMPS Technical Memo 1, Department of Computer Science, University of Utah, July 1980.
- [KL85] R.M. Keller and G. Lindstrom. Approaching distributed database implementations through functional programming concepts. In *Int'l Conference on Distributed Systems*, May 1985.
- [Kle36] S.C. Kleene.  $\lambda$ -definability and recursiveness. *Duke Mathematical Journal*, 2:340-353, 1936.
- [KR35] S.C. Kleene and J.B. Rosser. The inconsistency of certain forms of logic. *Annals of Mathematics*, 2(36):630-636, 1935.
- [Kro87] H.J. Kroeze. *The TWENTEL System (Version 1)*. Technical Report, Department of Computer Science, University of Twente, The Netherlands, 1986/87.
- [KS86] R.M. Keller and R. Sleep. Applicative caching. *ACM Transactions on Programming Languages and Systems*, 8(1):88-108, January 1986.
- [Lan64] P.J. Landin. The mechanical evaluation of expressions. *Computer Journal*, 6(4):308-320, January 1964.
- [Lan65] P.J. Landin. A correspondance between ALGOL 60 and Church's lambda notation. *Communications of the ACM*, 8:89-101,158-165, 1965.
- [Lan66] P.J. Landin. The next 700 programming languages. *Communications of the ACM*, 9(3):157-166, 1966.
- [LG88] J.M. Lucassen and D.K. Gifford. Polymorphic effect systems. In *Proceedings of 15th ACM Symposium on Principles of Programming Languages*, pages 47-57, January 1988.

- [MAGD83] J. McGraw, S. Allan, J. Glauert, and I. Dobes. *SISAL: Streams and Iteration in a Single-Assignment Language, Language Reference Manual*. Technical Report, Lawrence Livermore National Laboratory, M-146, July 1983.
- [Mar51] A.A. Markov. Teoriya algorifmov (Theory of algorithms). *Trudy Mat. Inst. Steklov*, 38:176-189, 1951.
- [McC60] J. McCarthy. Recursive functions of symbolic expressions and their computation by machine, Part I. *CACM*, 3(4):184-195, April 1960.
- [McC63] J. McCarthy. A basis for a mathematical theory of computation. In *Computer Programming and Formal Systems*, pages 33-70, North Holland, Amsterdam, 1963.
- [McC78] J. McCarthy. History of Lisp. In *Preprints of Proceedings of ACM SIGPLAN History of Programming Languages Conference*, pages 217-223, 1978. Published as SIGPLAN Notices 13(8), August 1978.
- [Mcg82] J.R. McGraw. The VAL language: description and analysis. *TOPLAS*, 4(1):44-82, January 1982.
- [Mil78] R.A. Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17(3):348-375, December 1978.
- [Mil84] R. Milner. A proposal for Standard ML. In *Proceedings 1984 ACM Conference on LISP and Functional Programming*, pages 184-197, ACM, August 1984.
- [MS76] R.E. Milne and C. Strachey. *A Theory of Programming Language Semantics*. Chapman and Hall, London, and John Wiley, New York, 1976.
- [Mul88] L.R. Mullin. *A Mathematics of Arrays*. PhD thesis, Computer and Information Science and CASE Center, Syracuse University, December 1988.
- [NPA86] R.S. Nikhil, K. Pingali, and Arvind. *Id Nouveau*. Computation Structures Group Memo 265, Massachusetts Institute of Technology, Laboratory for Computer Science, July 1986.
- [Pfe88] F. Pfenning. Partial polymorphic type inference and higher-order unification. In *Proceedings 1988 ACM Conference on Lisp and Functional Programming*, pages 153-163, ACM SIGPLAN/SIGACT/SIGART, Salt Lake City, Utah, August 1988.
- [PJ87] S.L. Peyton Jones. *The Implementation of Functional Programming Languages*. Prentice-Hall International, Englewood Cliffs, NJ, 1987.

- [PJCSH87] S.L. Peyton Jones, C. Clack, J. Salkild, and M. Hardie. GRIP—a high-performance architecture for parallel graph reduction. In *Proceedings of 1987 Functional Programming Languages and Computer Architecture Conference*, pages 98–112, Springer-Verlag LNCS 274, September 1987.
- [Pos43] E.L. Post. Formal reductions of the general combinatorial decision problem. *American Journal of Mathematics*, 65:197–215, 1943.
- [RCe86] J. Rees and W. Clinger (eds.). The revised<sup>3</sup> report on the algorithmic language Scheme. *SIGPLAN Notices*, 21(12):37–79, December 1986.
- [Rey74] J.C. Reynolds. Towards a theory of type structure. In *Proc. Colloque sur la Programmation*, pages 408–425, Springer-Verlag LNCS 19, 1974.
- [Rey85] J.C. Reynolds. Three approaches to type structure. In *Mathematical Foundations of Software Development*, pages 97–138, Springer-Verlag LNCS 185, March 1985.
- [Ros82] J.B. Rosser. Highlights of the history of the lambda-calculus. In *Proceedings 1982 ACM Conference on LISP and Functional Programming*, pages 216–225, ACM, August 1982.
- [Sch24] M. Schönfinkel. Über die bausteine der mathematischen logik. *Mathematische Annalen*, 92:305, 1924.
- [Sch85] D.A. Schmidt. Detecting global variables in denotational specifications. *ACM Transactions on Programming Languages and Systems*, 7(2):299–310, 1985.
- [Sco70] D.S. Scott. *Outline of a Mathematical Theory of Computation*. Programming Research Group PRG-2, Oxford University, November 1970.
- [SH86] Jr. Steele, Guy L. and W. Daniel Hillis. Connection machine lisp: fine-grained parallel symbolic processing. In *Proceedings 1986 ACM Conference on Lisp and Functional Programming*, pages 279–297, ACM SIGPLAN/SIGACT/SIGART, Cambridge, Massachusetts, August 1986.
- [Sha84] E. Shapiro. *Systolic Programming: A Paradigm of Parallel Processing*. Department of Applied Mathematics CS84-21, The Weizmann Institute of Science, August 1984.
- [Sri85] N.S. Sridharan. *Semi-Applicative Programming: An Example*. Technical Report, BBN Laboratories, November 1985.
- [Sto77] J.E. Stoy. *Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory*. The MIT Press, Cambridge, Mass., 1977.
- [Sto85] W. Stoye. *A New Scheme for Writing Functional Operating Systems*. Technical Report 56, University of Cambridge, Computer Laboratory, 1985.

- [TBH82] P.C. Treleaven, D.R. Brownbridge, and R.P. Hopkins. Data-driven and demand-driven computer architectures. *Computing Surveys*, 14(1):93–143, March 1982.
- [Tha87] S.S. Thakkar, editor. *Selected Reprints on Dataflow and Reduction Architectures*. The Computer Society Press, Washington, DC, 1987.
- [Tof88] M Tofte. *Operational Semantics and Polymorphic Type Inference*. PhD thesis, University of Edinburgh, Department of Computer Science (CST-52-88), May 1988.
- [TP86] H-C. Tu and A.J. Perlis. FAC: a functional APL language. *IEEE Software*, 3(1):36–45, January 1986.
- [Tra88] B.A. Trakhtenbrot. *Comparing the Church and Turing Approaches: Two Prophetic Messages*. Technical Report 98/88, Eskenasy Institute of Computer Science, Tel-Aviv University, April 1988.
- [Tu86] H-C. Tu. *FAC: Functional Array Calculator and its Application to APL and Functional Programming*. PhD thesis, Yale University, Department of Computer Science, April 1986. Available as Research Report YALEU/DCS/RR-468.
- [Tur36] A.M. Turing. On computable numbers with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42:230–265, 1936.
- [Tur37] A.M. Turing. Computability and  $\lambda$ -definability. *Journal of Symbolic Logic*, 2:153–163, 1937.
- [Tur76] D.A. Turner. *SASL language manual*. Technical Report, University of St. Andrews, 1976.
- [Tur79] D.A. Turner. A new implementation technique for applicative languages. *Software — Practice and Experience*, 9:31–49, 1979.
- [Tur81] D.A. Turner. The semantic elegance of applicative languages. In *Proceedings of the 1981 Conference on Functional Programming Languages and Computer Architecture*, pages 85–92, ACM, 1981.
- [Tur82] D.A. Turner. Recursion equations as a programming language. In *Functional Programming and Its Applications: An Advanced Course*, pages 1–28, Cambridge University Press, 1982.
- [Tur85] D.A. Turner. Miranda: a non-strict functional language with polymorphic types. In *Functional Programming Languages and Computer Architecture*, pages 1–16, Springer-Verlag LNCS 201, September 1985.

- [Veg84] S.R. Vegdahl. A survey of proposed architectures for the execution of functional languages. *IEEE Transactions on Computers*, C-23(12):1050–1071, December 1984.
- [vH67] J. van Heijenoort. *From Frege to Gödel*. Harvard University Press, Cambridge, MA, 1967.
- [Vui74] J. Vuillemin. Correct and optimal implementations of recursion in a simple programming language. *Journal of Computer and Systems Science*, 9(3), 1974.
- [WA85] W.W. Wadge and E.A. Ashcroft. *Lucid, the Dataflow Programming Language*. Academic Press, London, 1985.
- [Wad] P. Wadler. Efficient compilation of pattern-matching. Chapter 5 in [PJ87].
- [Wad71] C.P. Wadsworth. *Semantics and Pragmatics of the Lambda Calculus*. PhD thesis, Oxford University, 1971.
- [Wad86] P. Wadler. A new array operation. In *Workshop on Graph Reduction Techniques*, Springer-Verlag LNCS 279, October 1986.
- [Wad87] P. Wadler. *Views: A Way for Pattern-Matching to Cohabit with Data Abstraction*. Technical Report 34, Programming Methodology Group, Chalmers University of Technology, March 1987. Preliminary version appeared in the Proceedings of the 14th ACM Symposium on Principles of Programming Languages, January 1987.
- [WB89] P. Wadler and S. Blott. How to make *ad hoc* polymorphism less *ad hoc*. In *Proceedings of 16th ACM Symposium on Principles of Programming Languages*, pages 60–76, January 1989.
- [Weg68] P. Wegner. *Programming Languages, Information Structures, and Machine Organization*. McGraw-Hill, 1968.
- [Wik88] Wikstrom. *Standard ML*. Prentice-Hall, 1988.
- [Wis87] D. Wise. Matrix algebra and applicative programming. In *Proceedings of 1987 Functional Programming Languages and Computer Architecture Conference*, pages 134–153, Springer Verlag LNCS 274, September 1987.
- [WM88] P. Wadler and Q. Miller. *An Introduction to Orwell*. Technical Report, Programming Research Group, Oxford University, 1988. (First version, 1985.).
- [WW87] P. Watson and I. Watson. Evaluating functional programs on the FLAGSHIP machine. In *Proceedings of 1987 Functional Programming Languages and Computer Architecture Conference*, pages 80–97, Springer-Verlag LNCS 274, September 1987.

[You88] J. Young. *The Semantic Analysis of Functional Programs: Theory and Practice*. PhD thesis, Yale University, Department of Computer Science, 1988.