

Fancy Ferns Require Little Care*
(Extended Abstract)
by

Daniel P. Friedman and David S. Wise†

Abstract: A fern is a kind of tree with at most one non-terminal node at every level. It is defined as a structure for encapsulating results of asynchronous, indeterminate multiprocessing. Formal properties of two structure builders, CONS and FRONS, are proved; most interesting are the idempotence properties regarding FRONS and 1, and the contrasting order-preserving behavior of CONS.

We present a simple language using these operations through axioms and examples. Finally, practical demands of shared references, programming practice, and fairness are demonstrated to refine the axiomatized language. This separation of reference-sharing issues from the definitions on indeterministic choice, and its subsequent derivation, are important steps toward formal semantics for indeterminism.

Keywords and phrases: asynchronous recursion, fern, amb, multiset, bag, frons, cons, list, prefix#, stream, indeterminism, call-by-need, parallel processing, suspension, lazy evaluation, unpredictable ordering, arbiter, arbit.

CR Categories: 4.20, 4.32, 4.34, 4.13, 4.35

*Headline The New York Times Sunday, Apr. 29, 1979, page D39
Vol. 128, Issue 44, 202.

†Research supported in part by The National Science Foundation under Grant numbers MCS77-22325 and MCS79-04183.

Introduction

This paper presents a formal development of ferns and integrates them into the environment of applicative programming practice. We have developed the operational semantics of ferns elsewhere [6, 5] to a degree far in advance of the theoretical thrust of this paper. Here we only put that work on as firm a footing as possible.

Ferns are defined as a data structure which may first be perceived as a generalization of list structures (a la [18][12]) designed to include an (initially) unordered structure here dubbed a multiset. Our perception of list is firmly rooted in an applicative-style programming, that precludes side-effects (assignment imperatives). We derive our generalization from the rule that a list's content, once the list is defined, is immutable. Since content need not even be evaluated at the time a list is constructed [2], a list built as a totally ordered object may include values represented as "suspensions," immutable information sufficient to generate those values as in call-by-need [21, 20, 7] semantics for parameter passing. Because these suspensions are analogous to dormant processes, parallelism can be introduced thereby [3].

A multiset [9] (also bag [22]) is essentially a list built without specifying explicit order. As we shall see in the last section of this paper, one can argue that order can be impressed upon the structure during its use, but for now we shall allow its order to remain unspecified except for one convention. As an unordered structure, its access is not deterministic and involves choice. We constrain that choice through mechanisms to be associated

With EUREKA images of e-abstraction, so that alternatives involving suspensions which are ill-behaved are avoided. In terms of denotational semantics [17], we will only allow choices associated with non-1 values (or with ordering forced at construction time - i.e., list orderings).

In introducing this convention, which we associate with the term "indeterminism", we argue briefly that call-by-need is a natural evaluation strategy for operational semantics. A critical ability of computers (aside from making decisions) is the ability to recall rapidly values already known. Typically this is done by fetching the value of a variable from memory; in lambda-calculus the analog is recovering the bound value of a variable. In lambda-calculus models various protocols are available to effect this during the evaluation of an expression: call-by-name, call-by-value, call-by-need; all have more or less the same effect in terms of Church-Rosser equivalence. The last two, however, are operationally closest to the "table-look-up" behavior of computer memory efficient in time since no argument is evaluated twice. As we argue next, this avoidance of "twiceness" is necessary to grapple with indeterminism.

Indeterminism is a facility often used by a programmer to incorporate external behavior into his program through selection of "good" values and avoidance of "bad" ones. Such a facility specifies a choice which is not determined just by program or by input values, but may involve time dependencies among various input streams. (It is not to be confused with the non-determinism

of automata theory, a stylistic convenience for compressing explosive deterministic computations.) A program which uses indeterminism, having made an underspecified choice may refer repeatedly to the results of that choice.

We present a formal definition of ferns for representing indeterminism. The first section includes some notation, definitions, and examples of ferns - the generalization of lists but built from multisets. The definitions of CONS - the list builder, FRONS - the multiset-builder, and EUREKA - the choice builder provide for building and probing ferns. Formal results at this level establish the mathematical properties of these definitions that later support the language built on them. A congruence relation is established on the set of ferns and forms the basis for several important results: that FRONS is insensitive to 1 (Corollary 6.3 on Idempotency), that FRONS avoids only 1 (Corollary 7.1), and in contrast that CONS preserves order regardless of whether or not its domain argument is 1 (Theorem 8). Infinite ferns are also considered. Theorem 10, establishing that any choice available in decomposing an (unordered) fern is also available if that fern were COMSED in the first place, forms the foundation of the later language definition: that unordered ferns may be accessed as if they had been constructed specifying some total order (as CONS does).

This supports the semantics presented in the following

section, which provides that a fern appears to be any one of its CONS • EUREKA images (where • denotes composition), the choice of image being made on every application of the ε-abstraction. There the user is provided a mutating frons, μfrons, constructor in addition to McCarthy's cons [12]. Examples there include amb, arbit, symmetric or, and MERGE of arbitrary numbers of streams. In the final section we justify the restriction of μfrons to a non-mutating frons, available without changing any axioms. As a result, an indeterministic choice need only be made once within the language and preserved within a fern structure.

Conventional Notation

A sequence over an alphabet is any string composed of symbols from the alphabet. The empty sequence is Λ . If V is an alphabet, then V^* denotes the set of all finite sequences over V , including Λ . V^∞ denotes the set of all infinite sequences over V . A sequence S may be perceived as a partial function from the natural numbers, ω , to the alphabet. Thus we might write $S \in V^*$ as $S = s_0 \cdot s_1 \cdot s_2 \dots$ for $s_i \in V$. Because there are symbols in our alphabet which are not represented by a single character we require the syntactic use of a concatenation symbol "." between symbols in a sequence. Some alphabets include sets as elements.

We also use the conventional powerset operator:

$PS = \{X \mid X \subseteq S\}$ and function extension operator:

$f[y/x] = \lambda u.(u=x + y, fu)$ where the $(+,)$ form is the strict conditional. Finally, we will be referring to an arbitrary Scott Domain, D , and its bottom element (\perp) .

A multiset over D is a mapping, $M: D \rightarrow \omega+1$. We interpret the multiset as a function mapping $d \in D$ into the number of times (perhaps infinite) that it occurs in the multiset. The set of multisets over D is $M(D)$ abbreviated to M where obvious. Singleton multisets are written using set notation; if $d \in D$ then $\{d\} = (\lambda x.x=d + 1, 0) \in M(D)$. We define additive operators on multisets analogous to set union and difference [9 22].

If $M, N \in M(D)$ then

$$M \oplus N = \lambda x.M(x) + N(x);$$

$$M \ominus N = \lambda x.M(x) - N(x).$$

Definition: The set of ferns over D is

$$F(D) = \{A_0 \cdot a_1 \cdot A_2 \cdot a_3 \cdot A_4 \cdot \dots \in (M(D) \cdot D)^* \cdot M(D) \cup (M(D) \cdot D)^\infty \mid \forall i \geq 0, A_{2i+1} (a_{2i+1}) > 0\}.$$

We write an element of

$F(D)$ in Old English font: $A \in F(D)$. If $A \in F(D)$ then we write $A = A_0 \cdot a_1 \cdot A_2 \cdot a_3 \cdot \dots$ where $A_{2i} \in M(D)$ and $a_{2i+1} \in D$. This indexing is taken to be conventional so that the meaning of A_{2i} and a_{2i+1} are all notationally implied once A is known. Similarly B_0 and b_1 (C_0 and c_1) when B (respectively C) is defined.

A trivial element of F is the empty fern $Nil = \lambda x.0$.

Example of a simple fern

We next present a simple example of three related ferns. That these ferns are built from one another reflects the operational nature of the sharing inherent in an implementation.

The choice of the word "fern" (and later of "FRONS") is motivated by the shape of these structures as illustrated (Figure 1) in the following examples. The reader is encouraged to compare those figures with that of equisetum in Figure 2, which is not quite a tree.

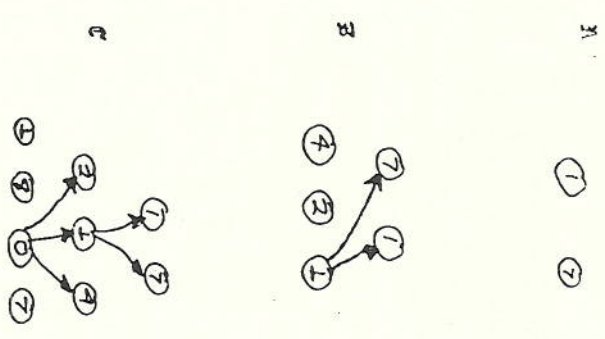



Figure 1. The example ferns A, B, and C.



FLORA OF INDIANA
Owen
Chas. C. Deam, collector

Figure Z. Equisetum arvense (Courtesy of Herbarium of Indiana University).

$$A = A_0$$

$$A_0 = \{1, 7\}$$

$$A = B_0 \cdot b_1 \cdot B_2$$

$$B_0 = \{2, 4, 1\}$$

$$b_1 = 1$$

$$B_2 = A_0$$

$$C = C_0 \cdot c_1 \cdot C_2 \cdot c_3 \cdot C_4$$

$$C_0 = \{1, 8, 0, 7\}$$

$$c_1 = 0$$

$$C_2 = B_0$$

$$c_3 = b_1$$

$$C_4 = B_2 = A_0$$

A list is a fern, $K \in F$, such that K_{21} is a singleton or $K_{21} = \lambda x.0$ for all x . Note that a list is composed of singleton multisets. For example, the LISP [12] list (1 2 1 3) is represented by $X = \{1\} \cdot 1 \cdot \{2\} \cdot 2 \cdot \{1\} \cdot 1 \cdot \{3\} \cdot 3 \cdot \lambda x.0$.

The purpose of CONS and FRONS is to build up ferns as similar constructors build up lists in LISP [12]. The function EUREKA specifies a set of possible pairs which specify how a fern is to be probed in a "bottom-avoiding" fashion. It is reminiscent of other choice functions: amb [12], arbit [8], and merge [1] which avoid results of ill-defined or misbehaved computation, but we use it only descriptively: it is not part of a user's language.

Notation: For the following three definitions we specify a local string notation on ϕ and ψ . Define ϕ and ψ by $A = A_0 \cdot \phi$ and, if $\phi \neq \Lambda$ then $\phi = a_1 \cdot A_2 \cdot \psi$. This notation references the suffices after the first one or three items in A .

CONS: $D \times F \rightarrow F$

$(d, A) \mapsto (d) \cdot d \cdot A$

FRONS: $D \times F \rightarrow F$

$(d, A) \mapsto (d) \otimes A_0 \cdot \phi$

EUREKA: $F \rightarrow P(D \times F)$

$A \mapsto \{ \langle a_1, (A_2 \otimes A_0 \otimes \{a_1\}) \cdot \psi \rangle \mid A_0 = \{a_1\} \vee a_1 \neq 1 \}$
 $\cup \{ \langle a, (A_0 \otimes \{a\}) \cdot \phi \rangle \mid a \neq 1 \ \& \ A_0(a) > 0 \ \& \ (a = a_1) \Rightarrow (A_0(a) > 1) \}$

One way of forming our example fern is by generating

\mathfrak{H} from A and then \mathcal{C} from \mathfrak{H} :

$A = \text{FRONS}(1, \text{FRONS}(7, \text{Nil}))$

$\mathfrak{H} = \text{FRONS}(2, \text{FRONS}(4, \text{CONS}(1, A)))$

$\mathcal{C} = \text{FRONS}(8, \text{FRONS}(1, \text{FRONS}(7, \text{CONS}(0, \mathfrak{H}))))$.

EUREKA for ferns Λ , \mathfrak{H} , and \mathcal{C} follows:

$\text{EUREKA}(A) = \{ \langle 1, (7) \rangle, \langle 7, (1) \rangle \}$

$\text{EUREKA}(\mathfrak{H}) = \{ \langle 2, \{1, 4\} \cdot 1 \cdot A \rangle, \langle 4, \{1, 2\} \cdot 1 \cdot A \rangle \}$

$\text{EUREKA}(\mathcal{C}) = \{ \langle 0, \{1, 1, 2, 4, 8, 7\} \cdot 1 \cdot A \rangle, \langle 7, \{1, 0, 8\} \cdot 0 \cdot \mathfrak{H} \rangle, \langle 8, \{1, 0, 7\} \cdot 0 \cdot \mathfrak{H} \rangle \}$.

Theorem 1: $\text{FRONS}(d, \text{Nil}) = \{d\}$; $\text{CONS}(d, \text{Nil}) = \{d\} \cdot d \cdot (\lambda x. 0)$.

Definition: If $A \in F$ then a promotion sequence of F is a sequence $\{ \langle d_i, A_i \rangle \}_{i > 0} \in (D \times F)^{\omega} \cup (D \times F)^{\infty}$ such that $A_0 = A$ and $\langle d_{i+1}, A_{i+1} \rangle \in \text{EUREKA}(A_i)$ for $i \geq 0$. A promotion sequence of length n is maximal if either $n < \omega$ and $\text{EUREKA}(A_n) = \emptyset$ or $n = \omega$.

There are infinite ferns whose maximal promotion sequences are all finite: for example, $\lambda x. 1$ whose triviality anticipates Corollary 3.2 and Corollary 6.3.

One (of several) promotion sequences of \mathfrak{H} is:

$\langle 4, \mathfrak{H}_1 \rangle, \langle 2, \mathfrak{H}_2 \rangle, \langle 1, \mathfrak{H}_3 \rangle, \langle 7, \mathfrak{H}_4 \rangle, \langle 1, \mathfrak{H}_5 \rangle$ where

$\mathfrak{H}_1 = \{1, 2\} \cdot 1 \cdot A$

$\mathfrak{H}_2 = \{1\} \cdot A$

$\mathfrak{H}_3 = A$;

$\mathfrak{H}_4 = \{1\}$

$\mathfrak{H}_5 = \text{Nil}$.

Definition: We define the binary relation $<$ on F by

$A < B$ when for every maximal promotion sequence $\{ \langle d_i, A_i \rangle \}_{i > 0}$ (of A) there is a promotion sequence $\{ \langle d_i, B_i \rangle \}_{i > 0}$ (of B),

The promotion sequence of B need not be maximal, but may be longer than that of A so long as the $\{d_i\}$ coincide for the length of A 's maximal promotion sequence.

Theorem 2: $A < \mathbb{N}$ if and only if for every $\langle d, A' \rangle \in \text{EUREKA}(A)$ there exists $\langle d, \mathbb{N}' \rangle \in \text{EUREKA}(\mathbb{N})$ such that $A' < \mathbb{N}'$.

Corollary 2.1: The $<$ relation is reflexive and transitive.

Definition: The binary relation $=$ on F is defined by $A = \mathbb{N}$ if and only if $A < \mathbb{N}$ and $\mathbb{N} < A$.

Corollary 3.1: The $=$ relation is an equivalence relation on F .

Using the fern $\mathbb{N} = B_0 \cdot b_1 \cdot B_2$ from our earlier example, form $\mathbb{N}' = \text{FRONS}(1, \mathbb{N}) = \{1\} \circ B_0 \cdot b_1 \cdot B_2$. The additional 1 in B_0 interferes with the construction of maximal promotion sequences for \mathbb{N}' equal in length to those for \mathbb{N} . In particular for each promotion sequence $\langle d_1, \mathbb{N}' \rangle_{i>0}$ for \mathbb{N}' there is no $d_1' = 1$; whereas $d_3 = 1$ in every maximal promotion sequence $\langle d_1, \mathbb{N} \rangle_{i>0}$ for \mathbb{N} . Hence $\mathbb{N}' < \mathbb{N}$ but $\mathbb{N}' \neq \mathbb{N}$. On the other hand consider $\mathbb{N}'' = B_0 \cdot b_1 \cdot B_2 \circ \{1\}$. In this case there is no such interference and the maximal promotion sequences of \mathbb{N}'' and \mathbb{N} are the same; hence $\mathbb{N}'' = \mathbb{N}$.

The next set of theorems explores what happens to promotion sequences as a fern is built up using FRONS and CONS. It turns out that if $A = \text{FRONS}(1, \text{Nil})$ then $\text{EUREKA}(A) = \emptyset$ and A has only empty promotion sequences; hence $A < \text{Nil}$. As we expect, EUREKA avoids 1 and so the promotion sequences don't change with 1 in the fern.

Corollary 3.2: If $A < \text{Nil}$ then $A = \text{Nil}$.

In general, however, $\text{FRONS}(1, A) \neq A$, but a weaker relationship can be established. We explore this fact in the general case of FRONS and later study CONS.

Theorem 4: If $A < \mathbb{N}$ then $\text{FRONS}(d, A) < \text{FRONS}(d, \mathbb{N})$.

Now we can justify the use of the congruence symbol, \equiv .

Corollary 4.1: If $A = \mathbb{N}$ then $\text{FRONS}(d, A) = \text{FRONS}(d, \mathbb{N})$.

Therefore, \equiv acts like a left congruence. More conventionally we would choose an infix "binary" operator, $\circ = \text{FRONS}$, as the operator preserving this congruence.[†]

Theorem 5: $\text{FRONS}(1, A) < A$.

We are tempted by the discussion preceding Corollary 3.2 to establish 1 as an analogous "left identity".

That would follow from a result symmetric to Theorem 5, which is not available. Unfortunately, promotion sequences of A can be disturbed by "FRONSing" a 1: Consider $A = \text{FRONS}(1, \text{FRONS}(1, \text{Nil}))$.

$\mathbb{N} = \text{CONS}(1, \text{Nil})$ has the promotion sequence $\langle 1, \text{Nil} \rangle$ of length one, but $\text{EUREKA}(A) = \emptyset$. Thus $A < \mathbb{N}$ but not $\mathbb{N} < A$. We can, however, establish an idempotence result somewhat analogous to Theorem 5.

Lemma 6.1: Let $A \in F$ then if $\forall i \ a_{2i+1} \neq 1$ then $A < \text{FRONS}(1, A)$.

[†] An equivalence relation \equiv is a left congruence with respect to an operator \circ if $x \equiv y$ implies $z \circ x \equiv z \circ y$.

Lemma 6.2: Let $A \in F$ and k be the smallest integer

such that $a_{2k+1} = 1$. If $k = \omega$ or if $A_{2k}(1) > 1$ or if

there exists $j < k$ such that $A_{2j}(1) > 0$, then $A < \text{FRONS}(1, A)$.

We now have our idempotence results.

Theorem 6: $\text{FRONS}(1, A) < \text{FRONS}(1, \text{FRONS}(1, A))$.

Corollary 6.3: $\text{FRONS}(1, A) = \text{FRONS}(1, \text{FRONS}(1, A))$.

Corollary 6.4: $\lambda y. \text{FRONS}(1, y)$ is idempotent with respect to $=$.

Thus, FRONSing 1 once is the same as FRONSing it repeatedly.

Theorem 7: If $d \neq 1$ then $\text{FRONS}(d, A) < A$ iff

$\exists k \omega (\forall j < k \ a_{2j+1} = d \ \& \ (\exists A_{2i}(d) = \omega \text{ where } 0 \leq i \leq k))$.

Corollary 7.1: Let $\mathbb{N} = \text{FRONS}(d, A)$.

Then $\mathbb{N} < A$ iff $\mathbb{N} = A$ or $d = 1$.

Now that we have thoroughly explained the behavior of FRONS under the interpretation above, we present similar results about the more familiar CONS operation. We, however, lose the idempotence.

Lemma 8.1: $\text{EUREKA}(\text{CONS}(d, A)) = \{ \langle d, A \rangle \}$.

Theorem 8: If $\text{CONS}(d, A) < A$ then $(\exists A_{2i}(d) = \omega \text{ where } 0 \leq i \leq \omega)$.

Corollary 8.2: Let $\mathbb{N} = \text{CONS}(d, A)$.

Then $\mathbb{N} < A$ iff $\mathbb{N} = A$.

Nevertheless, $=$ is not a left congruence under CONS.

Theorem 9: If $A = \mathbb{N}$ then $\text{CONS}(d, A) = \text{CONS}(d, \mathbb{N})$.

Theorem 10: If $\langle d, \mathbb{N} \rangle \in \text{EUREKA}(A)$ then $\text{CONS}(d, \mathbb{N}) < A$.

Semantics

In this section we present a semantics for a language that includes ferns as data structure for the user. Most of the axioms are conventional -- lifted from Stoy [19]. Ferns are embedded in an unconventional way, however, outside that tradition.

Two perspectives, operationally motivated, are necessary to understand why we refrain from using traditional power domain approaches to indeterminism. First is the call-by-need [21] or call-by-delayed-value [20] approach to parameter passing and to data structure construction [2]. We intend that no expression ever be evaluated until necessary to further the fabrication of an outermost result. In delaying evaluation we include specifically the content of ferns and of environments (treated also as objects -- albeit function-like). Such content, therefore, may be shared through borrowed references before it is ever evaluated. Upon first access such postponed evaluation proceeds at most once. There is no loss in such an operational protocol compared to that of conventional substitution for deterministic programs; we are considering more.

The second perspective, then, is the necessity to assure consistent handling of twice-used results of indeterminism. A major contribution of this paper is the two-step approach to this requirement. In the axioms which follow, indeterminism is introduced via Axiom 11 (FRONS) and is used via Axiom 12. In a later section we use these same axioms to reach a more comfortable solution to "twiceness", embedding all choice within ferns.

Axioms

- $\mathcal{E}[I]_p = \text{oid}$ where $I \in \text{Id}$ (1) Identifier
- $\mathcal{E}[\text{true}]_p = \text{tt}$ (2) Truth
- $\mathcal{E}[\text{false}]_p = \text{ff}$ (3) Falseness
- $\mathcal{E}[\text{nil}]_p = \text{Nil}$ (4) Empty fern
- $\mathcal{E}[E_0 = E_1]_p = (\mathcal{E}[E_0]_p = \mathcal{E}[E_1]_p)$ (5) Equality
- $\mathcal{E}[\text{strictify}(E_0, E_1)]_p = (\mathcal{E}[E_0]_p = 1 + 1, \mathcal{E}[E_1]_p)$ (6) Sequencer

The strictify operator is a sequencing operator, which is used to control the convergence properties of indeterminate structures.

We also specify a conditional†

$$\mathcal{E}[\text{if } E_0 \text{ then } E_1 \text{ else } E_2]_p = \mathcal{E}[E_0]_p + \mathcal{E}[E_1]_p, \mathcal{E}[E_2]_p \quad (7) \text{ Conditional}$$

Function invocation is usually conventional, although

Axiom 12 is related in an unusual way.

$$\mathcal{E}[\lambda x. E]_p = \lambda x. \mathcal{E}[E]_p[x/I] \quad (8) \lambda\text{-abstraction}$$

$$\mathcal{E}[E_0 E_1]_p = \mathcal{E}[E_0]_p \mathcal{E}[E_1]_p \quad (9) \text{ Application}$$

† More cleanly we might have defined an additive conditional [19] and then defined "if E_0 then E_1 else E_2 " = "strictify(E_0 , addf $E_0 E_1 E_2$)" but we do not need additive conditionals in this paper [4].

$\mathcal{E}[\text{cons}(E_0, E_1)]_p = \text{CONS}(\mathcal{E}[E_0]_p, \mathcal{E}[E_1]_p)$ (10) Determinate Constructor

We use μfrons locally in the user language to specify a fern which appears to mutate under Axiom 12. In the next section we define another user primitive frons from μfrons.

$\mathcal{E}[\text{μfrons}(E_0, E_1)]_p = \text{FRONS}(\mathcal{E}[E_0]_p, \mathcal{E}[E_1]_p)$ (11) Indeterminate Constructor

While cons and μfrons are themselves always deterministic, they build ferns which can only be accessed using Axiom 12.

$\mathcal{E}[\text{eI}.E]_p = \lambda A. (\mathcal{E}[E]_p[d/I, \mathcal{R}/J])$ where $\langle d, \mathcal{R} \rangle \in \text{EUREKA}(A)$ (12) e-abstraction

Axiom 12 uses the call-by-need perspective on environment -- that arguments are only evaluated once -- to bind two variables consistently. The e-abstraction borrowed from λ-abstraction and "εupnκ" specifies that one choice of a pair is made on each application, and that these two items simultaneously and consistently become the values to which two variables are bound.

Delayed evaluation (call-by-need) may postpone this simultaneous binding, but it is not possible to use call-by-name which might separate it. For example, a call-by-name protocol might allow

$\mathcal{E}[\text{eI}. \text{μfrons}(I, J)]_p$

where $\mathcal{E}[F]_p = \mathcal{F} = \{tt, ff\}$

to yield $\{ff, ff\}$

if I were bound to ff and J were bound to

$\{ff\}$ independently. (But

$\text{EUREKA}(\mathcal{F}) = \{\langle tt, \{ff\} \rangle, \langle ff, \{tt\} \rangle\}$.)

Axiom 12 does not allow the substitution to distribute the choice across any other operator. Another way to perceive the situation is that the choice is made using call-by-need in the "environment" built on application of an e-abstraction.

Thus, the conventional power domain constructions [16, 18] do not work here. The EUREKA-image of a fern, not in general effectively computable, is only used as the domain of a choice function in Axiom 12. There, however, two bindings are effected in a necessarily consistent fashion, precluding call-by-name.

Examples

The examples in this section fall into three categories:

simple, recursive on D , and recursive on $D + F$. The simple functions are amb [12] and arbit [8]. The recursive on D examples are the infinite sequence of integers, in; a permutation of the integers; sequential or[13,pg58] symmetric or[14,pg46]. The recursive on $D + F$ example merge streams.

The use of recursion in the programming language deserves some justification. Since we have not established monotonicity or continuity results we cannot use least-fixed-points to guarantee a meaning for recursion. Instances of λ -abstractions are in the language, however, so from a syntactic perspective we can construct Curry's Y-combinator and apply it to functions to admit recursive examples [19 , p.73] (solving no equations) under normal order evaluation, which is provided through call-by-need.

This scheme allows us to specify the construction of infinite ferns, but never to actually construct one. As long as A_0 is finite for any fern A , moreover, we can use a scheduling strategy to compute the elements of A_0 in parallel in order to find one member of EUREKA(A). If A_0 is infinite, then the search for an element of EUREKA(A), for a $\neq 1$, is more complicated; a scheduling strategy akin to the conventional enumeration of the rational numbers suffices.

We first present fern probing relations:

```
first = e!j.i;
rest = e!j.j.
```

Since first(F) and rest(F) can have as many possible values as the cardinality of EUREKA(F) and these can be chosen independently on each application, it is difficult to use these probes on ferns built directly by ufcons. We, therefore, don't use first' and rest here; in the next section this restriction will be relaxed.

The simple functions are non-recursive and include amb and arbit. Amb is a function which takes two arguments and returns a non-1 one of them if it exists.

```
amb = \A.\B.(exy.x) ufcons(A,ufcons(B,nil)) .
```

Arbit is similar to amb, but it returns tt if the first one is non-1, false if the second one is non-1, and 1 otherwise.

```
arbit = \A.\B. amb strictly(A,true) strictly(B,false)
```

Of the functions that are recursive on D we have the natural numbers, in, and the sequential and symmetric or.

```
in = \I.C(I,n(I+1))
```

If C is cons then we have an infinite sequence; if C is ufcons each of its promotion sequences is a permutation of the natural numbers.

or = $\lambda F. \text{if null } F \text{ then false}$

$\text{else } (\text{eUV. if } U \text{ then true else or } V) F$.

If a list is bound to F then the behavior is the sequential or; if a multiset is bound to F then the behavior is the symmetric or. Or has meaning for any fern, not just lists and multisets.

We next present one piece of code, which

may be interpreted as four different programs. One of these is the "interrupt handler" or "input driver" required to service any active terminal in the airline reservation problem [1, 23,].

Before presenting these examples we extend D, hitherto an arbitrary domain, to include ferns.[†] This provides for ferns of ferns, sublists, and other nested data structures as anticipated by McCarthy [12].

Consider the code

$\text{merge} = \lambda M. \text{if null } M \text{ then nil}$

$\text{else } (\text{eLR. if null } L \text{ then merge } R$

$\text{else } (\text{eUV. } C(U, \text{merge } X(V,R))) L) M.$

where C and K are the constructors ufrons or cons.

[†] Showing that F is a domain is beyond the scope of this paper.

This code is quite similar in effect to code for flatten when $C=\text{cons}=K$. Such a function reduces a list of lists (or a matrix) into a single list of its second level elements (or a vector). When the argument is less well ordered more interesting things happen.

Let us assume that each sublist of L (or row of the matrix) has been constructed with Landin's prefix# [11], a version of cons which is strict in its first argument. Such sublists, or streams, are characteristic of serial input lines in a communication network. If these streams are composed into a multiset of substreams using ufrons, then we have the analog of a matrix with order and "seriality" with each row, but whose rows occur in an unspecified order.

The effect of $C=\text{cons}=K$ is then to append all streams, the streams to be chosen in an order depending on the (temporal) order of convergence of the first elements in the streams. Once a stream is chosen to be first, however, it must be exhausted before another one is chosen to be merged in.

The effect of $C=\text{cons}$ and $K=\text{ufrons}$ is the desired interrupt handler. Upon each recursion the multiset structure is restored so that the output is a true merging of convergent prefixes of active streams. The order of the merge is determined.

If $C = \mu\text{frons} = K$, the result is a multiset of all elements in all streams without regard to order either within streams or among streams. The stream nature of input, however, tends to assure that internal order within a stream will be reflected in the shuffled output, but this observation is based solely on operational behavior rather than required semantics.

If $C = \mu\text{frons}$ and $K = \text{cons}$ then the shuffling of the previous paragraph is restricted up to the first infinite stream encountered. C assures that the result is an unordered structure, but K forces a selected stream to be exhausted before admitting the next one into the shuffle.

Coping with anomalous behavior -- frons

The programmer working with indeterminism and the language of the previous section has sufficient power to solve some interesting problems, but he has some irritation as well. For example, let $\mathcal{E}[F]_p = \mathcal{F}$ be a fern and let EUREKA(\mathcal{F}) have two elements $\langle d, \mathcal{F} \rangle$ and $\langle d', \mathcal{F}' \rangle$. Now within the scope of one application of ϵ , say

$$\mathcal{E}[\epsilon I J . (I = I) E]_p$$

the values of I and J are consistently paired, so that this evaluation yields \underline{tt} always: either I is bound to d or to d' exclusively. However, every application of ϵ allows a new choice for the pair:

$$\mathcal{E}[(\epsilon I J . I E) = (\epsilon I J . I E)]_p$$

may evaluate either to \underline{ff} or to \underline{tt} depending on whether d and d' are chosen by either application of ϵ or not. That the equality of this example might not hold should be quite disturbing to the applicative programmer because the syntax is identical on either side of the equal sign.

The headaches compound if one grapples with "fairness."

A fair implementation of Axiom 12 would have to select an element from a fern's EUREKA image with probability equal to all the rest. Consider the infinite multiset of natural numbers; a possible result of $\mathcal{E}[\lambda F . F \text{ nn}(0)]_p$:

$$\mathcal{E}[F]_p = \mathcal{N} = \lambda x . i \in F(\omega).$$

Then every fair evaluation

$$\mathcal{E}[(\epsilon I J . I) F]_p$$

should turn up "yet another" integer; this should act like a random Integer Generator* which is not effectively computable. Fairness is, therefore, hopeless.

* This observation about fairness and multisets is due to Robert E. Filman.

Recalling our early discussion of "twiceness" and call-by-need, our programmer perceives that the problem lies in repeated applications of Axiom 12 to the same fern, yielding (legally) inconsistent results and precluding fairness. Theorem 10 shows a way out: fix it so that the application of ϵ is embedded within a fern. Then the restrictions on evaluating arguments once will apply Axiom 12 only once per fern:

$\text{frons} = \lambda X.\lambda Y.(eif.\text{cons}(I,J)) \text{ufrons}(X,Y).$

By Theorem 10 $\llbracket \text{frons}(X)Y \rrbracket_p < \llbracket \text{ufrons}(X,Y) \rrbracket_p$ so that frons is a correct, but weak replacement for ufrons; its behavior is one of the possible behaviors of ufrons.

In effect, a fern built with frons in place of ufrons everywhere, is built as a list incorporating one of the EUREKA choices possible at each step of a promotion sequence. That list reflects one, not all, maximal promotion sequences. Of course, the call-by-need convention serves to delay selection of precisely which one; the choice unfolds simultaneously with the need for such a sequence.

The anomalous behavior of repeated applications of ϵ , seen to yield inequality between syntactic identicals before, now disappears. Every application of ϵ yields consistent unique results, as if the consistency rules for ϵ have moved outside its scope. (Indeed, this is just what the cons in the definition of frons does: it freezes the ϵ -scope in which it is evaluated.) frons is no more a function than is ufrons, but first and rest become deterministic functions in a style where frons displaces all uses of ufrons.

Moreover, the fairness of frons is an issue different from the fairness of ϵ . If frons replaces ufrons in all code (e.g. the examples of the previous section), then there is no longer any fairness debate on ϵ . All remaining occurrences of ϵ (aside from that in the definition of frons itself) are deterministic, so we need only concern ourselves with the fairness of frons (or that one occurrence of ϵ). An important refinement has been made if ϵ is only indeterminate within frons: we know that ϵ is never necessarily applied to the same fern, A , twice unless $A_0 = \{a\}$ a singleton.

Thus, a concern for fairness need not account for fair relational (as opposed to functional) behavior of Axiom 12 choosing from the EUREKA image of one fern repeatedly. Under frons only one indeterminate choice is ever made for every fern, so that fairness need never be defined over independent applications of Axiom 12; instead we can establish some dependency according to the order of the maximal promotion sequence selected. Just such a proposal was the original goal of our work [5 , appendix].

Conclusion

We have developed a theoretical perspective on indeterminism for applicative multiprogramming. There have been several other approaches to incorporating indeterminism in pure applicative languages. Of these, the earliest is McCarthy's amb operator [12] which returned one or the other of its two arguments, avoiding \perp as a value when possible. Bottom-avoidance is also the essence behind Kosinski's arbiter [10] and Keller's arbit [8]. None of these extends to consideration of fairness.

By embedding the indeterminism within a structure -- the fern -- we allow references to the contention to be passed as arguments, establishing shared references to the choice before it is made. Call-by-need protocol for application and for structure definition not only allows reference to an unmade (but specified) choice, but also guarantees normal order evaluation and allows us recursive definitions using the Y combinator,

Encapsulating indeterminism inside an object allows us to attack fairness of that object, rather than having to relate fairness to overall control structure. If choice is made fairly as an object is built up and torn down -- over a period of time -- then the system has a fair choice mechanism available to it. In the preceding section we only anticipated solutions to fairness; here we offer a conjectural definition of what a fair implementation of ferns should do.

Fairness Principle : An interpretation of cons and ufcons may be said to be fair if there exists $f: P \rightarrow \omega$, some monotonically increasing function with respect to $<$ on P and to the total order on ω , such that $f(A)$ is bounded by the length of all maximal promotion sequences of A and f meets the following requirement. When $\mathcal{E} \upharpoonright P \cap \mathcal{A} = A_0$ with maximal promotion sequence $\langle d_1, A_1 \rangle_{1 > 0}$ and $\mathcal{E} \upharpoonright \text{ufcons}(d, F) \upharpoonright P = \mathcal{F}_0$ with maximal promotion sequence $\langle e_1, \mathcal{F}_1 \rangle_{1 > 0}$ then either $d = \perp$ and $\forall i (e_i = d_i)$ or $d \neq \perp$ and $e_1 = d_1$ for $1 < f(A_0)$, $e_i(f(A_0)) = d_i$, and $e_{i+1} = d_i$ for $i > f(A_0)$.

The fairness principle conjectures a function which identifies the prefix of a promotion sequence which is not altered by ufconsing new elements into the fern. Thus, new values are inserted after such an (ever increasing) prefix. This guarantees that every element in a fern will occur "early" in a maximal promotion sequence regardless of how many items are adjoined to that fern with ufcons; eventually the new values will take their place after already extent elements. Eventually any non- \perp element occurs "first" as the fern is traversed using Axiom 12. (The conjectured function, f , need not be locally properly increasing, but it must be asymptotically increasing for fairness to hold.) We have implemented one such scheme [5] using timestamps (birthdays) to effect the function f .

In summary, we have defined ferns as a data structure encapsulating indeterminism by extending lists. We have proven several results regarding the mathematical properties of ferns, particularly relating to the idempotency resulting from adding a contending process that doesn't behave; semantics of list structures are retained, however. Following a semantics for a language incorporating ferns we presented a few examples; others are available elsewhere [6,5]. Having stated the case for ferns carrying indeterminism within applicative languages, we offered a simple refinement to solve the twiceness anomaly of other formulations of indeterminism in applicative languages. That twist allows us to confront issues of fairness, set forth just above.

The challenge of fairness is that contention of processes in an applicative system could neither be ignored nor honored without any explicit scheduler to assure such effect. It is this challenge which we expect ferns to meet. Future work will refine the concept of fairness and reconcile it within a generalized formulation of denotational semantics.

Acknowledgement: We thank Steven D. Johnson and particularly Mitchell Wand for numerous constructive discussion and critical readings. Early encouragement from John Backus and James H. Morris, Jr., pushed this work along. We also thank Carl Hewitt who suggested the airline reservation problem and Robert Tennent and David MacQueen who offered some thoughtful proposals on earlier versions of this paper.

REFERENCES

1. Dennis, J. B. A language design for structured concurrency. In J. H. Williams and D. A. Fisher (eds.), Design and Implementation of Programming Languages, Springer, Berlin (1977), 231-242.
2. Friedman, D. P., and Wise, D. S. CONS should not evaluate its arguments. In S. Michaelson and R. Milner (eds.), Automata, Languages and Programming, Edinburgh Univ. Press, Edinburgh (1976), 257-284.
3. Friedman, D. P., and Wise, D. S. Aspects of applicative programming for parallel processing. IEEE Trans. Comput. C-27, 1 (April, 1978), 289-296.
4. Friedman, D. P., and Wise, D. S. A note on conditional expressions. Comm. ACM 21, 11 (November, 1978), 931-933.
5. Friedman, D. P., and Wise, D. S. An approach to fair applicative multiprogramming. In G. Kahn and R. Milner (eds.), Proc. of Intl. Symp. on Semantics of Concurrent Computation, Springer, Berlin (1979), to appear.
6. Friedman, D. P., and Wise, D. S. An indeterminate constructor for applicative programming. Proc. 7th ACM Symp. on Principles of Programming Languages (1980), 245-250.
7. Henderson, P., and Morris, J. H., Jr. A lazy evaluator. Proc. 3rd. ACM Symp. on Principles of Programming Languages (1976), 95-103.
8. Keller, R. M., Lindstrom, G., and Patil, S. Data-flow concepts for hardware design. COMPCON Conf. (1980),
9. Knuth, D. E. The Art of Computer Programming 2, Semi-numerical Algorithms, Addison-Wesley, Reading, MA (1969), 551.
10. Kosinski, P. R. A data flow language for operating systems programming. Proc. ACM SIGPLAN-SIGOPS Interface Meeting, SIGPLAN Notices 8, 9 (September, 1973), 89-94.
11. Landin, P. J. A correspondence between ALGOL 60 and Church's lambda notation. Comm. ACM 8, 2 (February, 1965), 89-101.

12. McCarthy, J. A basis for a mathematical theory of computation. In P. Braffort and D. Hirschberg (eds.), Computer Programming and Formal Systems, North-Holland, Amsterdam (1963), 33-70.
13. McCarthy, J., Abrahams, P. W., Edwards, D. J., Hart, T. P., and Levin, H. E. LISP 1.5 Programmer's Manual, M.I.T. Press, Cambridge, MA (1962), Chapter 1.
14. Manna, Z. Mathematical Theory of Computation, McGraw-Hill, New York (1974), Chapter 5.
15. MacLane, S. Categories for the Working Mathematician, Springer-Verlag, New York (1971) 26.
16. Plotkin, G. D. A powerdomain construction. SIAM J. Comput. 2, 3 (September, 1976), 452-487.
17. Scott, D. S. Logic and programming languages. Comm. ACM 22, 9 (September, 1977), 634-641.
18. Smyth, M. B. Power domains. J. Comp. Sys. Sci. 16 (1978) 23-36.
19. Stoy, J. E. Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory, M.I.T. Press, Cambridge, MA (1977).
20. Vulliamin, J. Correct and optimal implementation of recursion in a simple programming language. J. Comp. Sys. Sci. 9, 3 (June, 1974), 332-354.
21. Wadsworth, C. Semantics and Pragmatics of Lambda-calculus. Ph.D. dissertation, Oxford (1971).
22. Waldinger, R. J., and Levitt, K. N. Reasoning about programs. Artificial Intelligence 5, 3 (Fall, 1974), 235-316.
23. Yonezawa, A., and Hewitt, C. E. Modelling distributed systems. 5th Intl. Joint Conf. on Artificial Intelligence (1977), 370-377.

*Laboratory for
Programming Methodology*

Symposium

on

Functional Languages

and

Computer Architecture

Invited Papers