

THE
UCLA
COMPUTER SCIENCE
DEPARTMENT
QUARTERLY

THE DEPARTMENT
AND ITS PEOPLE

FALL 1987/WINTER 1988

VOL. 16 NO. 1

COMPUTER SCIENCE DEPARTMENT OFFICERS

Dr. Gerald Estrin, Chair
Dr. Jack W. Carlyle, Vice Chair, Planning & Resources
Dr. Sheila Greibach, Vice Chair, Graduate Affairs
Dr. Lawrence McNamee, Vice Chair, Undergraduate Affairs
Mrs. Arlene C. Weber, Management Services Officer
Room 3713 Boelter Hall

QUARTERLY STAFF

Dr. Sheila Greibach, Editor
Ms. Marta Cervantes, Coordinator

**Material in this publication may not be reproduced without written permission of
the Computer Science Department.**

Subscriptions to this publication are available at
\$20.00/year (three issues) by contacting Ms. Brenda Ramsey,
UCLA Computer Science Department, 3713 Boelter Hall,
Los Angeles, CA 90024-1596.

**THE UCLA COMPUTER SCIENCE DEPARTMENT
QUARTERLY**

**The Computer Science Department
And Its People**

**Fall Quarter 1987/Winter Quarter 1988
Vol. 16, No. 1**

**SCHOOL OF ENGINEERING AND APPLIED SCIENCE
UNIVERSITY OF CALIFORNIA, LOS ANGELES**

Fall-Winter CSD Quarterly

CONTENTS

COMPUTER SCIENCE AT UCLA AND COMPUTER SCIENCE IN THE UNITED STATES

Gerald Estrin, Chair 1

THE UNIVERSITY CONTEXT 13

THE UCLA COMPUTER SCIENCE DEPARTMENT

The Faculty

Areas of interest	15
Biographies	18
Publications.....	33
Honors and awards	40

The Students

Biographies	41
Ph.D. Degrees Awarded In Calendar Year 1987	57
Ph.D. Students In Advanced Stage	58
M.S. Degrees Awarded In Calendar Year 1987	62
M.S. Students In Advanced Stage.....	64
The Staff	66
The Visitors	67
The Supporters	68

DEPARTMENT ORGANIZATION

Committees	70
Administrative Assignments	71
Activities.....	73
Student Organizations	76

ADMISSIONS DEADLINES, FACULTY POSITIONS

AND AFFIRMATIVE ACTION

78

COMPUTER SCIENCE DEPARTMENT TECHNICAL REPORTS

80

TECHNICAL PAPERS

On Style, Expressability, and Efficiency in Functional Programming Languages

Grabriel Robins

106

On Style, Expressibility, and Efficiency in Functional Programming Languages

Gabriel Robins

Computer Science Department
University of California, Los Angeles

Abstract

A functional style of programming was proposed by Backus as an alternative to conventional programming paradigms. Functional programming is intended to alleviate what Backus dubbed as the "Von Neumann Bottleneck," a restrictive set of biases deeply ingrained in the thought processes of programmers. The functional programming paradigm is investigated here through the actual implementation of a series of common algorithms. One of the novel features of this paradigm is the lack of side-effects; it is shown how this property can be used to achieve considerable efficiency improvement in arbitrary implementations of functional languages, using the idea of *value-caching*. This scheme is then shown to reduce the execution of certain computations from exponential-time to polynomial-time, a formidable speedup. It is also shown that the expressive power of functional languages can be greatly enhanced via the utilization of both *macros* and *idioms*, without changing the original semantics of the language.

1. Introduction

A functional style of programming was proposed in 1978 by John Backus [Backus] as an alternative to conventional programming paradigms. Functional programming is intended to alleviate what Backus dubbed as the "Von Neumann Bottleneck," a restrictive set of biases which by now is deeply ingrained in the thought processes of programmers. Backus argues quite eloquently that because the processor and the memory of a conventional computer are connected via a narrow bus, processors end up spending most of their time ferrying useless data to and from the main memory.

The more serious problem, however, is that the classical "word-at-a-time" model of computer design has over the years managed to permeate up through the abstraction levels and become fused into most programming languages, and indeed into the very thought processes of the majority of computer scientists. According to Backus, this phenomenon is directly responsible for many horrendously complex programming languages with convoluted and even inconsistent semantics, where it is often impossible to prove formally even the most trivial correctness properties.

Backus gave a syntax and semantics for a candidate functional programming language. Some of the novel features of his language are the lack of variable names, assignment statements, global state, and side-effects. Instead, several powerful combining operators can be used to compose functions, yielding clean and simple semantics, conciseness of expression, and easily-derivable formal correctness properties. See [Vegdahl] for a more detailed discussion of these issues. Considerable work regarding FP has also been performed at UCLA: [Alkalaj], [Alkalaj, Ercegovac, and Lang], [Arabe], [Meshkinpour], [Meshkinpour and Ercegovac], [Patel and Ercegovac], [Patel, Schlag, and Ercegovac], [Pungsornruk], [Sausville], [Schlag 86], [Schlag 84], [Worley].

In this paper I investigate some of the claims made by proponents of functional programming, through the actual implementation of a series of common algorithms. Based on these examples, I then draw some conclusions with regards to the practical advantages and flaws of the functional programming paradigm. The language I use is Berkeley's functional programming language, FP [Baden]. All FP forms which appear in this paper were verified to be correct by actually typing them into the FP v. 4.2 (4/28/83) system. The rest of the paper assumes that the reader is familiar with Berkeley FP; for a brief refresher on the syntax and semantics of FP, the reader is referred to the appendix.

I show how the lack of side-effects can be used to achieve considerable efficiency improvement in arbitrary implementations of functional languages, using the idea of *value-caching*. This scheme is proved to reduce the execution time of certain FP computations from exponential to polynomial, a remarkable speedup. Finally, it is shown that the expressive power of functional languages can be greatly enhanced via the utilization of both *macros* and *idioms*, without changing the original semantics of the language.

2. FP Solutions to Some Common Problems

In this section we examine how various common problems and algorithms can be handled in FP. We do this in order to get some feeling as to what does "typical" FP code look like. The notation " $X \implies Y$ " will be used to denote that when the form X is evaluated it yields the value Y , as in the following: "+ @ [%1 , %2] \implies 3".

2.1. Permutations

The following function produces a permutation of the second argument, according to the order specified by the first argument. That is, given the following data vector of length $N: < d_1, d_2, \dots, d_N >$, and a permutation vector $< p_1, p_2, \dots, p_N >$, where $1 \leq p_i \leq N$ and $p_i \neq p_j$ if $i \neq j$, we would like to return the vector $< d_{p_1}, d_{p_2}, \dots, d_{p_N} >$. The following FP function will accomplish this task:

```
{PERM1 &pick @ distr}
```

For example:

```
PERM1 : <<3 1 2 5 4> <is going what here on>> \implies <what is going on here>
```

But now one might argue that any straight-forward implementation of FP would cause the "distr" operator to "copy" its second argument a number of times proportional to the length of the first argument, using $O(N^2)$ space in total. In attempting a response to such criticism we may modify our definition as follows:

```
{PERM2 (null @ 1 -> [] ; apndl @ [pick @ [1 @ 1, 2], PERM2 @ [tl @ 1, 2]])}
```

```
PERM2 : <<2 4 3 1> <permuted 4 get things>> \implies <4 things get permuted>
```

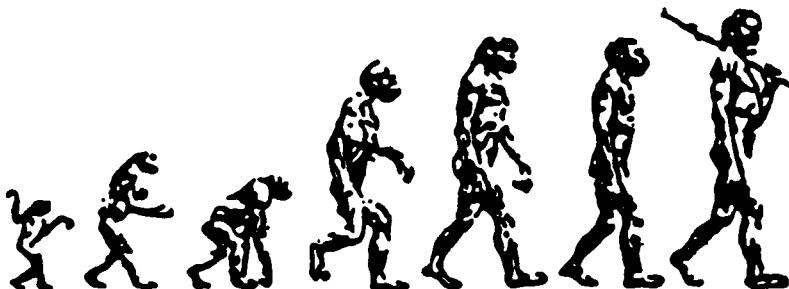
Now we are using tail-recursion to accomplish the same task. In constructing the seemingly "more space-efficient" function PERM2, we were presumptuous in assuming that the underlying FP implementation allocates space and executes FP forms in a certain way, and regardless of how close to the truth our assumptions may have been, these are the very kinds of considerations Backus is trying to free us from in the first place.

In programming in FP we must give up many familiar and useful concepts, such as variable names, a state, side-effects, etc. But we gain many pleasant properties in return for our "sacrifice": conciseness, more easily provable correctness, better parallelism, and freedom from details. If during FP programming we worry about efficiency and implementation issues, the efficiency of our code will become implementation-dependent, and our minds will be shackled by the very details from which we escaped when we sought refuge from the "Von Neumann bottleneck" in the first place.

If we are to enjoy the full benefits of a functional programming system, we must delegate most of the efficiency and implementation issues to the underlying (optimizing) compilers. One of the main benefits of FP is that it has very clean semantics, yielding relatively painless correctness proofs, and it is amenable to simple and powerful algebraic manipulations and transformations. Such algebraic transformations could be used to automatically convert/compile an FP program into a version which would take advantage of the inherent parallelism or other special features of the architecture we happen to be executing on.

In summary, I am *not* arguing that efficiency considerations are not important; I am simply

arguing that it is more the responsibility of the machines to catch up with functional paradigms of thinking and not the burden of functional programming to accommodate archaic architectures. The situation is dramatized in the following diagram:



Functional Programming



Hardware Support

2.2. List Reversal

Suppose we wanted to reverse a list. What first comes to mind is a recursive definition for the reversal operation, namely that the reversal of a list is equal to the last element followed by the reversal of the rest of the list. In FP this could be expressed as follows:

```
{REV1 (null -> id; apndl @ [last,REV1 @ tlr])}
```

For example:

```
REV1 @ iota : 13      ==>      <13 12 11 10 9 8 7 6 5 4 3 2 1>
```

In trying to determine whether any parallelism is inherent in list reversal, we try to apply the highly-parallel "tree-insert" FP operator, as in the following:

```
{REV2 | [2,1]}
```

```
REV2 @ iota : 13      ==>      <<<13 <12 11>> <10 <9 8>>> <<7 <6 5>> <<4 3>> <2 1>>>>
```

We observe that this simple function is almost what we want, except for the extra levels of list depth, so we revise our definition as follows:

```
{REV3 | (concat @ [(atom @ 2 -> [2] ; 2) , (atom @ 1 -> [1] ; 1)] )}
```

```
REV3 @ iota : 13      ==>      <13 12 11 10 9 8 7 6 5 4 3 2 1>
```

Now we have obtained a highly parallel, yet simple and concisely-expressed function to reverse a list. I consider this to be functional programming at its best.

2.3. Membership in a List

Suppose we needed a function that given an element and a list, would determine whether that element appears in that list. Our first attempt is the following:

```
{MEMB ! or @ &= @ distl}  
MEMB :<13 <3 6 4 13 98 -45 13 73>> ==> T
```

The idea here is to distribute the search key over the entire list and "or" together the results of all the comparisons of these pairs. Now suppose we know ahead of time that our list was sorted. Using the classical binary-search technique, we could speed up the computation to "require" at most $O(\log N)$ time by recursively searching half-lists, as follows:

```
{BS (null @ 3 -> = @ [1,1 @ 2];  
<= @ [1,1@rotr@2] -> BS @ [1,2]; BS @ [1,3])) @ apndl @ [1,split @ 2]}
```

Note that although this "optimization" appears to reduce the evaluation time, this may not really happen in reality. For example, a close examination of the code for BS would reveal that we have assumed that the operation "split" operates in constant time. If indeed split is implemented in such a way as to require time linear in the size of its argument, BS may in fact run slower than its seemingly "less efficient" counterpart MEMB.

This example again illustrates our previous argument that FP programmers should not try to "second-guess" the implementation, because they may loose. It is conceivable that in the future intelligent compilers could algebraically manipulate such expressions and, for example, replace a linear search with a binary search whenever it finds a monotonicity property to be preserved. So again, I claim that such "optimizations" should be left to the compilers.

Of course one must draw the line somewhere. Given that the problem of whether two programs are equivalent or not is undecidable, in the general case it is not possible to perform these optimizations automatically. Yet even undecidability results do not mean that special classes of instances could not be handled automatically, as well as efficiently. For example, it is probably too much to hope that a compiler would recognize an $O(N^5)$ -time max-flow algorithm and replace it with an equivalent but more efficient $O(N^3)$ time algorithm; on the other hand it is not far-fetched to expect that an intelligent compiler could determine that a certain piece of code behaves like a look-up table and consequently implement it as a hash-table rather than a list, yielding better efficiency.

2.4. Fibonacci Numbers

The Fibonacci numbers are defined as the sequence of integers beginning with two 1's with each subsequent integer being the sum of the previous two. This recursive definition immediately yields an FP program:

```
{F (<=@[id,%2] -> %1; +@[F @ D,F @ D @ D])} # returns the ith Fibonacci number  
{D -@[id,%1]} # decrements by 1  
{FIB &F @ iota} # returns the first n Fibonacci numbers  
  
FIB : 18 ==> <1 1 2 3 5 8 13 2134 55 89 144 233 377 610 987 1597 2584>
```

Although this program is both elegant and concise, it is quite "inefficient", because any conventional implementation of it would force the re-computation of the function on many of the same arguments. We can alleviate this flaw as follows:

```
{FIB2 tlr @ tlr @ FF @ [iota ,[%1,%1]]}  
{FF (null @ 1 -> 2 ; FF @ [tl@1 ,apndr @ [2, + @ [last@2 ,last@tlr@2]]])}
```

Now successive Fibonacci numbers are appended to a growing list, where the last two elements

are used directly to compute the next Fibonacci number. This optimization is essentially an algorithmic one and reduces the "execution" time from exponential-time to polynomial-time, a considerable savings. I believe such programmer-affected optimization to be quite valid, and also one which FP programmers should endeavor to concern themselves with. Later in this paper, we describe a method of achieving considerable running-time improvement while preserving the functional appearance of functions such as the Fibonacci numbers function; this is accomplished via the idea of value-caching. In other words, sometimes we can "have the cake and eat it too," functionally speaking.

2.5. Sorting

Sorting a list entails rearranging its elements in such a way as to preserve some monotonicity property with respect to successive elements. For simplicity, we assume that our elements are integers and that the monotone property we wish to sort with respect to is the normal arithmetic "less-than" or "<". Our first sorting algorithm is fashioned after the classical *selection-sort*:

```
{SS (null -> [] ; apndl @ [1, SS @ SUB] @ [MIN,id])}
{SUB concat @ &(<-> [] ; [2]) @ distl}
{MIN | (<-> 1 ; 2)}

MIN : <3 5 0 9 213 -7 1>      ==>      -7
SUB : <13 <2 13 a 5 w 13>>    ==>      <2 a 5 w>
SS : <4 7 5 1 0 9 11 7>        ==>      <0 1 4 5 7 9 11>
```

The function MIN determines the least element in a given list, while the function SUB returns a new list without any occurrences of a given element. Sorting is accomplished by finding the least element, putting it first, followed by the sorted list of the remaining elements.

Next we try another method for sorting, which is similar to our selection-sort algorithm. To sort a list, we keep rotating it until the first element is least. We then continue this process on the tail of the list, until the entire list is sorted. We call this algorithm *rotation-sort*:

```
{RS (null -> []; (= @ [1,MIN] -> apndl @ [1, RS @ t1]; RS @ rotr))}
{MIN | (<-> 1 ; 2)}
RS : <5 3 13 -2 -13 13 0 99 213> ==> <-13 -2 0 3 5 13 13 99 213>
```

As an interesting aside, our list is now allowed to have duplicates and such duplicate elements are not eliminated from the list. Although our rotation-sort algorithm is specified very concisely, it does not exhibit much parallelism and would require "quadratic" time to run in any conventional implementation. Another method of sorting entails swapping adjacent pairs of elements until the list becomes monotonic. We call this scheme *odd-even-sort*:

```
{SORT (while UNORD (apndl @ [1,SWAP @ t1] @ SWAP))}
{UNORD lor @ &> @ trans @ [tlr,t1]}
{SWAP concat @ &FLIP @ pair}
{FLIP (null @ t1 -> id; (>-> [2,1]; id))}

UNORD : <-5 2 5 9>          ==> F
FLIP : <13 2>                ==> <2 13>
SWAP : <76 13 0 12 33 22 0 13>    ==> <13 76 0 12 22 33 0 13>
SORT : <76 13 -12 22 0 13 213 0 >    ==> <-12 0 0 13 13 22 76 213>
```

The idea here is to look at the pair of elements 1 and 2, 3 and 4, etc., and swap each pair in-place if they are out of order. During the second pass, we similarly treat the pairs 2 and 3, 4 and 5, etc. We keep repeating this two-phased pass over the data until the list becomes sorted. Note the high degree of inherent parallelism here, since all the pairs may be examined/swapped in parallel.

A common method for sorting with a good average-time behavior is quicksort. The quicksort algorithm selects a pivot element and splits the input into two sublists, one consisting of elements less than (or equal-to) the pivot, and the other list consisting of elements greater-than the pivot

element. Each of these two sublists is similarly sorted in a recursive fashion, and the final result then consists of the first list, followed by the pivot element, followed by the second list. The FP implementation of *quicksort* follows:

```
{QS (<= @ [length,%1] -> id; concat @ [QS @ L, [1], QS @ R])}
{L concat @ & (>          -> tl; []) @ distl @ [1,tl]}
{R concat @ & (<=> tl; []) @ distl @ [1,tl]}

L : <7 3 9 -4 13 27>           ==> <3 -4>
R : <7 3 9 -4 13 27>           ==> <9 13 27>
QS : <76 13 -12 22 0 13 55 213 0 >    ==> <-12 0 0 13 13 22 55 76 213>
```

Again, we note the conciseness and elegance of the FP specification. Moreover, it is well-known that quicksort has an $O(N \log N)$ average-time behavior (although the normal worst-case behavior is still quadratic). Later in this paper we describe how the FP definition of quicksort can be made even more concise using macros.

2.6. Prime Numbers

A prime number is a positive integer with *exactly* two distinct divisors: itself and 1. We would like to specify a functional program which given an integer N , returns a complete list of primes between 2 and N , inclusively. Here is our first pass at this problem:

```
{PRIMES concat @ &(PRIME -> [id] ; []) @ tl @ iota}
{PRIME land @ &(not @ = @ [%0,id]) @ &mod @ distl @ [id,tl@iota@ -@ [id,%1]]}

PRIMES : 20      ==> <2 3 5 7 11 13 17 19>
```

The idea here is to process all the integers between 2 and N , trying to determine which ones have no factors other than themselves and 1. All integers satisfying this property are primes and are therefore returned as a list. The function PRIME is a predicate which tests a single integer for primality, while the function PRIMES maps PRIME onto an entire list of candidates for primes.

We note however, that to test for primality, we need only try to divide by factors no greater than the square root of our candidate integer, since if our integer is divisible by a factor which is greater than the square root, the result of the division is less than the square root and so we would have tried that integer earlier. Furthermore, other than the integer 2, only odd integers need be considered for primality, and only odd integers (up to the square-root of the candidate) need be considered as factors, since an even integer can not divide an odd one. These "optimizations" are reflected in a new version of our prime number generator:

```
{PRIMES apndl @ [%2, concat @ &(PRIME -> [id] ; []) @ ODDS]
{PRIME land @ &(not @ = @ [%0,id]) @ &mod @ distl @ [id,ODDS @ SQRT]}
{SQRT 1 @ (while (< @ [*@[id,id]@1, 2]) [+@[1,%1],2] @ [%1,id])
{ODDS apndl@[%2,&(+@[*@[id,%2],%1])
@iota@-@[+@[@[id,%2],mod@[id,%2]],%1]]}
```

```
PRIMES : 50      ==> <2 3 5 7 11 13 17 19 23 29 31 37 41 43 47>
```

Note that while this scheme is "faster" than the previous one, the code is beginning to look rather ugly. In fact, further speedup is possible since when checking for the primality of a candidate integer, we need only to consider other primes (up to the square-root of the candidate) as possible factors. The lesson here again is that the more "efficient" a functional program is, the messier it tends to become.

In the case of prime numbers, for example, I think that only the first "optimization" (the one about going only up to the square root) is of the type with which the functional programmers need concern themselves, because it is "algorithmic" in nature, and relies on a non-trivial number-theoretic fact; it so happens that it is also the only one which increases the "efficiency" of our

functional program by a non-constant factor.

We give another algorithm for prime generation, the *sieve of Erastostenes*:

```
{SIEVE concat @ &(MEMB -> []|[1]) @ distr @  
    [id,concat @ &&* @ &distl @ distr @ [id,id]] @ tl @ iota)  
(MEMB !or @ &= @ distl)  
  
SIEVE : 18      ==> <2 3 5 7 11 13 17>
```

The idea here is to list all the integers and erase all factors of 2, all factors of 3, and so on. Everything which remains after this process of elimination is by necessity a prime. Actually, in this implementation we do the opposite: we form all possible products and then consider as prime all the "missing" integers. Again we see an example of the trade-off between concise specification and computational "efficiency".

3. FP vs. APL

There are some notable similarities between APL and FP. FP is quite similar in spirit and syntax to APL. Some functions and operators in FP are taken directly from APL; for example, the FP iota function is equivalent to the APL iota function ("t"). Similarly, the FP operator "right-insert" ("!") is equivalent to the APL operator "reduce" (""). The difference with respect to these operators is that in FP one can compose any number of operators together in a concise and semantically clean manner, while in APL only a pre-defined limited set of types of compositions are allowed.

Both in FP and APL complex operations may be defined very concisely, often giving rise to "one-line" programs. Yet unlike FP, the semantics of APL are very complex, being muddled by state, local and global variables, parameter passing, operator overloading, goto statements, and the existence of several hundreds of different functions and operators. I have seen one-line APL programs which take the better part of an hour to understand; a possible moral here is that "absolute power corrupts absolutely." While it is true that some pieces of FP code may also seem rather dense, the situation in FP is still much better than in the notorious APL "one-liners".

4. FP vs. LISP

FP is also similar in many respects to LISP. Both languages rely heavily on lists as their primary data structures, and both languages provide a rich set of list-manipulation primitives. On the other hand, LISP functions are more defined than their FP counterparts; for example in FP `tl : <>` is undefined, while in LISP (`CDR NIL`) is still NIL. Given that FP functions are bottom-preserving, it is often the case that an FP function would return "?" on a boundary or pathological datum, while the LISP equivalent would still return the "right" answer.

Here is an example of when it is preferable to have functions which are more defined on "boundary conditions". Given an atom `x` and a list of atoms `y`, we would like to return the set of all atoms in `y` which directly follow occurrences of `x` in `y`. The obvious way of specifying this in FP elegantly is:

```
{F concat @ &(@ [2,1@1] -> [2@1] ; []) @ distr @ [trans @ [tlr,tl] @ 2 , 1]}  
  
F : <3 <m 3 4 5 6 3 3 x 3>>      ==>      <4 3 x>  
F : <y <>>                      ==>      ?
```

The problem here of course is that our function is not defined for the pathological case when the second argument is the empty list, while it would be preferable to define the answer in that case to be the empty list. We could modify our original function to check for that boundary condition as a special case, but if the FP functions `tl` and `tlr` were originally designed to return the empty list when called upon the empty list, our function `F` would not have the problem mentioned here. For example, if `tl`, `tlr`, and `trans` were redefined thus:

```
{TL (= @ [ length , %0 ] -> [] ; tl )}
```

```
{TLR (= @ [ length , %0 ] -> [] ; tlr ))
{TRANS (land @ &null -> [] ; trans)}
```

F could now use these "more defined" specification to work correctly in all cases, including the boundary conditions:

```
{F concat @ &(@ [2,1@1] -> [2@1] ; []) @ distr @ [TRANS @ [TLR,TL] @ 2 , 1]}
```

```
F : <y <x>>    ==>    <>
F : <y >>    ==>    <>
```

The point we wish to stress is that had some of the primitive functions been just a little more "defined", the programmer would not have had to worry about these boundary conditions, and freeing the programmer from such details is clearly a goal of the functional-thinking paradigm.

5. Macros in FP

FP currently does not allow the use of macros. This is, however, an artificial restriction since macros can be specified and used in a way that will not alter the original semantics of FP whatsoever. FP macros specify simple "static" substitution rules that would "expand" certain pieces of FP code. In other words, a macro-expansion preprocessor would be invoked before execution ever commences.

To accommodate macros, the original FP syntax would be extended as follows:

```
macro-definition ==> '#define' name '(' argument-names ')' macro-body
name          ==> letter (letter | digit | '_')*
argument-names ==> name ( ',' name)*
macro-body     ==> (function-form | name)*
function-form  ==> simple-function | composition | construction
                  | conditional | constant | insertion | alpha
                  | while | '(' function-form ')' | macro-expansion
macro-expansion ==> name '(' arguments ')'
arguments      ==> string ( ',' string)*
string         ==> (letter | number | special-character)*
```

It is understood that a macro expansion form and the corresponding macro-definition must agree on the number of arguments. It is further observed that arguments to macros may include arbitrary strings, not just FP objects, as long as a valid FP form is obtained when macro-expansion is terminated. Moreover, macro-expansion is a recursive process, where macros may refer to other macros, as long as termination of the macro-expansion process is guaranteed.

As an example, the macro `#define addl(x) + @ [x,%1]` would expand the form `addl(k)` into the form `+ @ [k,%1]`. As a second example of using macros, our previous definition of quicksort could be specified more concisely as:

```
#define S(X) concat @ &(X -> tl; []) @ distl @ [1,tl]
{QS (<= @ [length,%1] -> id; concat @ [QS @ S(>=), [1], QS @ S(<)])}
```

The above form, after the macro-expansion into "pure" FP, will become:

```
{QS (<= @ [length,%1] -> id; concat @ [QS @
                           concat @ &(>= -> tl; []) @ distl @ [1,tl], [1], QS @
                           concat @ &(< -> tl; []) @ distl @ [1,tl]])}
```

Note that in this example, we can not achieve the same effect by simply defining a new function, because the argument we really need to pass is a functor itself, not a value. In general, macros will increase the readability of FP code by making the FP representation of a function more concise, whenever certain pieces of code are identical or at least very similar. A macro can then be used to "factor out" that common piece of code, so that it does not have to be specified or

expressed more than once; in particular, macros would allow the definition of an APPLY template, similar in spirit to the LISP apply function, something that is impossible to implement in "pure" FP.

5.1. The Implementation of Macros

We emphasize that the usage of macros as defined here does not extend nor change the original semantics of FP in any way. In fact, an implementation of this idea might consist of a macro-expansion pre-processor which would take as input a specification in macro-enhanced FP (that is, an FP program with some macros) and output an equivalent "pure" FP program after properly doing the requisite macro expansion and substitution. The result would then be fed into the original FP processor.

For example, to implement a macro pre-processor for Berkeley FP, one could extend the original FP syntax to accommodate macros as described above, and construct a UNIX lex(1) and yacc(1) specification to lexically analyze and parse the new syntax. The associated production rules would perform the required macro expansion, producing as output an equivalent FP program in the original "pure" syntax. The function "yparse" can then be used as a "filter" to FP.

An easier alternative to using heavy-artillery such as yacc(1) and lex(1) consists of the following scheme: since many existing languages already have rather fancy macro-preprocessors, why not "fool" such a tool into doing our FP macro-processing work for us? For example, the standard UNIX C compiler has such a macro-preprocessor, and in fact it may be invoked without the actual C compilation. So an implementation of this scheme would consist of the following two steps:

- 1) Run the file containing "macro-enhanced" FP program through the C macro preprocessor and collect the resulting "pure" FP output in a file. This is accomplished by specifying the -E flag: `cc -E macro-prog.fp > pure.fp`
- 2) Run the resulting file as usual through the normal FP system: `fp < pure.fp`

I have already tested this scheme with several examples, including the quicksort example stated above, and it worked flawlessly. In fact, if one combines the above two steps into one UNIX command aliased to an appropriate name, the entire macro-expansion process would be transparent to the user. Of course, if the format of our proposed FP macros deviates at all from that of C, this trivial scheme would fail and we would be forced to revert back to our direct-translation scheme using lex(1) and yacc(1).

Macros could also be used to "modify" the syntax of FP by providing alternate ways of expressing FP constructs; such extensions to FP in no way alter the original semantics of FP but simply provide some "syntactic sugar" to ease the programming task.

6. The Need for Idioms

I found that during FP programming, many constructs repeat themselves. For example, the function to decrement a counter, to sort, etc. Such common and repeatedly used functions can and should be provided for in the implementation of an FP-like language. Indeed such functions can be easily specified within FP itself, thereby not adding to the original set of "primitives", but instead be provided as a *library* of functions or macros.

I claim that it would be very convenient to define certain idioms within FP. Note that this convenience is not in conflict with the fact that such idioms may be easily specified using other FP forms. For example, the original definition of FP still contains the idiom "length", even though "length" may easily be defined as a recursive application of `tl`, or perhaps even more cleverly as:

```
{length | + @ &%1}
```

Perlis and Rugaber have already devised such a set of idioms for APL [Perlis and Rugaber]. I believe that a similar set should be identified and defined for FP. In particular, I recommend the following list of possible idioms as a starting-point for the further development of this idea:

Removal of duplicates from a list - given a key x and a list y, return a new list z which is a copy of y except that all elements equal to x have been removed.

Sorting a list given a specified partial-order predicate - given a list, return the list sorted according to a specified (or default) comparison predicate. Since the said predicate is an argument to our "sort" routine, an implementation using macros would probably be best here.

Generalized 'iota' - given A, B, and C, return a sequence of all multiples of C, beginning with A and no larger than B.

Vector element uniqueness - given a sequence of objects, return "true" if there exists a repeated element in this sequence, or "false" otherwise.

Is X a permutation of Y? - given two sequences x and y, determine whether x is a permutation of y. This may be accomplished by sorting both into some "canonical order" and then comparing.

Minimum/maximum element - find the minimum (or maximum) element in a given sequence, using a specified min (or max) function, or else use a default comparison function.

Membership in a vector/set - test whether a given element x is a member of vector/set y.

Union of vectors/sets - given two vectors/sets x and y, return a vector/set of elements which are either in x or in y, without duplicates.

Intersection of vectors/sets - given two vectors/sets x and y, return a vector/set of elements which are both in x and in y, without duplicates.

Subtraction of vectors/sets - given two vectors/sets x and y, return a vector/set of elements which are in x but not in y, without duplicates.

Subset of a vector/set - given two vectors/sets x and y, determine whether each element of x is also an element of y.

Transitive closure - given a sequence of pairs or objects, return the transitive-closure of this set, interpreting the pairs as a specification of a binary relation.

Math functions - square root, exponentiation, factorials, binomial coefficients, successor, predecessor, etc.

Flatten a list - return a list of all atoms in a given list, disregarding the original level (or nesting-depth) of each.

7. Improved Efficiency Through Value-Cashing

We have observed that there exists a sharp trade-off in certain computations between conciseness-of-expression and "computational efficiency." For example, in the case of generating Fibonacci numbers, the elegant version of the algorithm appears as follows:

```
{F (<=@[id,%2] -> %1; +@[F @ D,F @ D @ D])} # returns the ith Fibonacci number  
{D - @[id,%1]} # decrements by 1
```

We also observed earlier that alternative ways of specifying the Fibonacci function yield more complex and difficult to read specifications. So the question now is how efficiently can the current specification be used to compute the ith Fibonacci number on a conventional architecture, and whether the situation can be improved. We settle the former part of the question first, and argue that a conventional implementation would require exponential time to compute the current Fibonacci numbers specification.

It is well-known that the Nth Fibonacci number is given by the expression:

$$F_n = [\phi^n - (1-\phi)^n]/(2\phi-1) \text{ where } \phi = (1+\sqrt{5})/2$$

The constant $\phi \approx 1.618033989$ is also known as the golden ratio, and when N approaches infinity, the term $(1-\phi)^n$ approaches zero and therefore F_n approaches $\phi^n/(2\phi-1)$ but since ϕ is a constant, we have $F_n = \Theta(\phi^n)$ or that F_n grows exponentially. Since a conventional implementation can not be expected to solve such recurrence relations in closed-form, it must use recursion to compute F_n . But the problem is that F will be recomputed many times at the same values, and since the value of F_n must be obtained from the lowest levels of the recursion via repeated additions of the constant 1, the total number of additions will be equal to the value of F_n , or proportional to ϕ^n . It follows that a naive computation of the Fibonacci sequence from the recursive specification will require exponential time.

Of course, in order to remain within polynomial-time, we could explicitly remember old values, but this is contrary to the spirit of the functional paradigm. We would like to combine the elegance of the concise recursive specification with the "efficiency" of the "optimized" version.

A hybrid can actually be achieved through what I call *value-caching*. The idea is to save argument/value pairs for certain difficult-to-compute functions, so that those values could be looked up in the future instead of having to be recomputed. We define a new FP operator $@\$$ (read compose-cache) to have the exact same semantics as the familiar FP operator $@$, except that after the composition is computed, the result is stored away in such a way that next time the same composition needs to be performed, the result is already known and can be used immediately without any recomputation or further delay. Clearly, the new composition operation so defined has the exact same semantics as the normal composition operator, and moreover, this entire scheme remains transparent to the programmer.

Because in FP there are no side-effects, composing a function f with an argument g will yield the exact same result at *any* stage of the computation, and that is why old results will *always* be valid in the future. This scheme can be implemented using hash-tables to store the value/result pairs for each function we would like to thus speed up. When a function is "called" upon an "argument", the argument is looked up in the hash-table associated with that function, and if a result is found it is returned immediately. If not, the result is computed from scratch, and then it is saved in the hash-table using the argument as a key, whereupon the result is passed on to the rest of the computation.

Of course not all functions should be so treated, only those that are non-trivial to compute. For example, using this scheme, our Fibonacci function would now appear as follows:

<pre>{F (<=@[id,%2] -> %1; +@[F @\\$ D,F @\\$ D @ D]))</pre>	# returns the i th Fibonacci num <pre>[D - @[id,%1]]</pre> # decrements by 1
--	--

Notice that we have not lost any of the conciseness or elegance apparent in our original specification. And assuming the cache/lookup process takes near-constant time, the Fibonacci function now will run in near-linear time on any conventional implementation so enhanced, an exponential improvement!

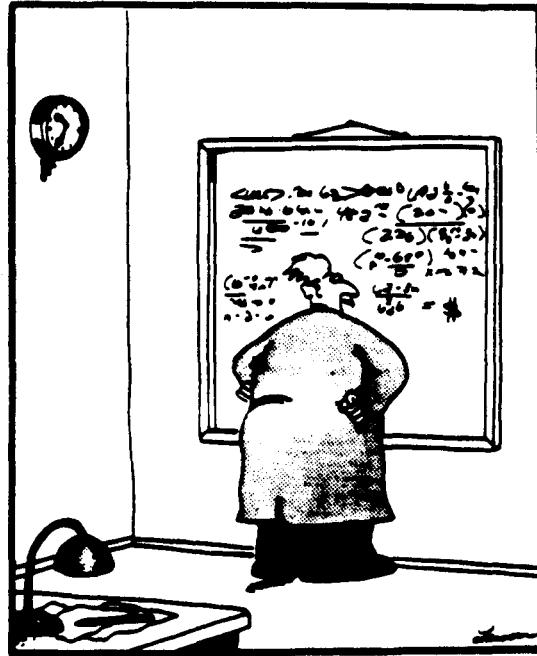
This scheme will work for arbitrary functional languages, not just FP. Moreover, the $@\$$ operator need not be used manually by the programmer, but simply be invoked automatically by the implementation whenever it is determined that doing so is likely to improve the efficiency. Some possible criteria for automatically performing value-caching are when considerable time has been spent inside some particular function that has relatively short input and output, when a particular function is being called a very large number of times on only a small set of different values, and when a function (or a set of functions) is highly recursive.

All of these situations can be detected automatically by the implementation, so it is possible to automate the idea of speed-up by value-caching. Sometimes, however, it may be the case that the programmer can better determine where caching should and should not be performed, so it may be desirable to give the programmer a "manual-override" capability with respect to value-caching. I propose two new variations on the composition operator, and a slightly modified behavior of the old composition operator:

@\$ - value-caching is *always* performed.

@^\$ - value-caching is *never* performed.

@ - architecture automatically decides when cashing should be performed.



Einstein discovers that time is actually money

Other instances where value-caching will yield considerable increase in execution efficiency are applications where computations have a geometric meaning in Euclidean D-space. For example, Conway's game of "life", where a cellular automata in the 2-dimensional plane is simulated and where the "next-generation" rule is a function of the local neighborhood surrounding the cell. A straight-forward recursive computation to determine the state of a given cell would run in time exponential in the magnitude of the generation to be computed, where the base corresponds to the size of the neighborhood around cells. On the other hand, the same algorithm running in a value-caching environment would run in polynomial time.

Generally speaking, whenever a dynamic-programming algorithm of some kind is used, new values are dependent on certain old values (with respect to some set of indices), and therefore an ordinary functional solution would likely to give rise to an inefficient execution. But if value-caching is used in the implementation, the execution time of the exact same algorithm is likely to improve drastically.

Recently it has come to my attention that the idea of caching values in order to avoid recomputations is also discussed in [Keller and Sleep]. They discuss several ways of storing and using old values in order to improve execution efficiency, and also address the problem of "purging" or keeping the cache size from growing without bounds during a computation. They also give as example another application area where value-caching would greatly speed up the execution time, namely in the finite-element method.

8. Conclusion

Through a series of examples I have investigated programming style and practical considerations

in FP. I have demonstrated that it is often the case in FP that concise functions become verbose and complicated when certain "efficiency optimizations" are attempted. I have also argued that certain so-called "optimizations" may really not be optimizations at all, and should not be undertaken by an FP programmer.

I have briefly discussed the similarities between FP and APL and LISP, and argued that certain FP functions should be "more defined" and that there is a need for some additional FP "primitive" functions or macros. I described a scheme where macros can be implemented in a clean way, leaving the original semantics of FP intact. I have argued that FP would benefit from a library of idioms, or commonly-used constructs upon which the programmer can draw.

It is possible via novel architectures and intelligent compilers to support efficient FP-like languages, where many optimizations can take place automatically, leaving the programmer free to see much of the high-level picture without drowning in the details. I introduced one such possible optimization, value-caching, and showed that it can reduce the execution time of certain computations by as much as an exponential amount.

9. Acknowledgments

I would like to thank Professor Milos Ercegovac for introducing me to FP, and for the numerous illuminating discussions which inspired this paper. Warm thanks go to Professor Sheila Greibach for encouraging me to polish up this paper and submit it to the quarterly; she made numerous thoughtful comments, and also helped in unsplitting my infinitives.

10. Bibliography

Alkalaj, L., "A Uniprocessor Implementation of FP Functional Language", MS Thesis, *UCLA Computer Science Department TR CSD-860064*, April 1986.

Alkalaj, L., Ercegovac, M., and Lang, T., "A Dynamic Memory Management Policy for FP", *1987 Hawaii International Conference on Systems Science*.

Arabe, J., "Compiler Considerations and Run-Time Storage Management for a Functional Programming System", Ph.D. Dissertation, *UCLA Computer Science Department, CSD TR 86004*, August 1986.

Backus, J., "Can Programming Be Liberated from the Von Neumann Style? A Functional Style and its Algebra of Programs", *Communications of the ACM*, Vol 21, No. 8, August, 1978, pp. 613-641.

Baden, S., *Berkeley FP User's Manual*, Rev. 4.1, September 29, 1987.

Keller, R., Sleep, R., "Applicative Caching: Programmer Control of Object Sharing and Lifetime in Distributed Implementations of Applicative Languages", *Proceedings of the 1981 ACM Conference on Functional Programming Languages and Computer Architectures*, October, 1981, pp. 131-140.

Meshkinpour, F., "On Specification and Design of Digital Systems Using an Applicative Hardware Description Language", M.S. Thesis, *UCLA Computer Science Department, Technical Report No. CSD-840046*, November 1984.

Meshkinpour, F., and Ercegovac, M., "A Functional Language for Description and Design of Digital Systems: Sequential Constructs", *IEEE Proc. of the 22nd ACM/IEEE Design Automation Conference*, 1985, pp.238-244.

Patel, D., and Ercegovac, M., "A High-Level Language for Hardware-Oriented Algorithm Design," presented at the *Symposium Electronics and Communications in the 80's*, Indian Institute of Technology, Bombay, February 1983.

Patel, D., Schlag, M., and Ercegovac, M., "vFP: An Environment for the Multi-Level

Specification, Analysis, and Synthesis of Hardware Algorithms," *Proc. 1985 ACM Conference on Functional Programming Languages and Architectures*, Nancy, France, Springer-Verlag Lecture Notes 201, 1985.

Perlis, A., Rugaber, S., "The APL Idiom List", *Research Report #87*, Yale University.

Pungsornruk, S., "A Threaded FP Interpreter/Compiler", M.S. Thesis, *UCLA Computer Science Department TR CSD-860061*, March 1986.

Sausville, M., "Gathering Performance Statistics on Hardware Specified in FP", *UCLA Computer Science Department Internal Report*, March 20, 1986.

Schlag, M., "Layout from a Topological Description", *Ph.D. Dissertation*, UCLA Computer Science Department, Summer 1986.

Schlag, M., "Extracting Geometry from FP Expressions for VLSI Layout", *UCLA Computer Science Report CSD-840043*, October 1984.

Vegdahl, S., "A survey of Proposed Architectures for the Execution of Functional Languages", *IEEE Transactions on Computers*, Vol c-33, No. 12, December, 1984, pp. 1050-1071.

Worley, J., "A Functional Style Description of Digital Systems", M.S. Thesis, *UCLA Computer Science Department, Technical Report CSD-860054*, February 1986.

10. Appendix: A Brief Overview of FP

In order to keep this paper self-contained, we briefly review the syntax and semantics of FP; for a more detailed description of FP, see [Baden].

The notation " $X \Rightarrow Y$ " is used to denote that when the form X is evaluated it yields the value Y ; this is simply a notational convenience. FP objects consist of numbers, strings, atoms, and lists; examples are %213, "hi there", FOO, <I am <a list>>, and <>. An integer constant must be preceded by a percent sign, to distinguish it from the primitive selector functions. We use the form " $F : x$ " to denote the result of the function F applied to the argument x ; for example, $+ : \%2 \%3 \Rightarrow 5$. The unique error atom (or "bottom") is denoted as "?". In what follows, the formal FP specifications are also explained informally for clarity.

$K : x \Rightarrow \text{if } x = <x_1, x_2, \dots, x_n> \text{ and } 1 \leq k \leq n \text{ then } x_k \text{ else ?}$

For every integer K there exists a primitive selector function K that returns the K^{th} element of a list. For example:

$3 : <\text{fi fy foo fum}> \Rightarrow \text{foo}$

Other selector primitives:

```
pick : x \Rightarrow \text{if } x = <y, x_1, x_2, \dots, x_n> \text{ and } 1 \leq y \leq n \text{ then } x_y \text{ else ?}
last : x \Rightarrow \text{if } x = <> \text{ then } <> \text{ else if } x = <x_1, x_2, \dots, x_n> \text{ and } 1 \leq n \text{ then } x_n \text{ else ?}
first : x \Rightarrow \text{if } x = <> \text{ then } <> \text{ else if } x = <x_1, x_2, \dots, x_n> \text{ and } 1 \leq n \text{ then } x_1 \text{ else ?}
tl : x \Rightarrow \text{if } x = <y> \text{ then } <> \text{ else if } x = <x_1, x_2, \dots, x_n> \text{ and } 2 \leq n \text{ then } <x_2, x_3, \dots, x_n>
else ?
tlr : x \Rightarrow \text{if } (x = <>) \text{ or } (x = <y>) \text{ then } <> \text{ else if } x = <x_1, x_2, \dots, x_n> \text{ and } 2 \leq n
\text{ then } <x_1, x_2, \dots, x_{n-1}> \text{ else ?}
```

`pick` selects the y^{th} element out of a given list, `last` selects the last element, `first` selects the first element, `tl` chops the first element off and returns the rest (like the LISP `cdr` function), and `tlr` chops the last element off and returns the rest. Examples are:

<code>pick : <2 <a b c>></code>	\Rightarrow	b
<code>last : <w x y z></code>	\Rightarrow	z

first : <<1 2> <3 4 5> <6 7 8 9>>	==>	<1 2>
tl : <who are you>	==>	<are you>
tlr : <fi fy fo fum>	==>	<fi fy fo>

Structural and list-manipulation primitives:

distl : x ==> if x=<y,>> then <>		
else if x=<y,<x ₁ ,x ₂ , ...,x _n >> and 1≤n then <<y,x ₁ >,<y,x ₂ >, ...,<y,x _n >> else ?		
distr : x ==> if x=<>,y> then <>		
else if x=<<x ₁ ,x ₂ , ...,x _n >,y> and 1≤n then <<x ₁ ,y>,<x ₂ ,y>, ...,<x _n ,y>> else ?		
apndl : x ==> if x=<y,>> then <y>		
else if x=<y,<x ₁ ,x ₂ , ...,x _n >> and 1≤n then <y,x ₁ ,x ₂ , ...,x _n > else ?		
apndr : x ==> if x=<>,y> then <y>		
else if x=<<x ₁ ,x ₂ , ...,x _n >,y> and 1≤n then <x ₁ ,x ₂ , ...,x _n ,y> else ?		
trans : x ==> if x=<<x ₁ ,x ₂ , ...,x _n >,<y ₁ ,y ₂ , ...,y _m >,...,<z ₁ ,z ₂ , ...,z _l >> and 1≤n		
then <<x ₁ ,y ₁ , ...,z ₁ >,...,<x _n ,y _n , ...,z _n >> else ?		
reverse : x ==> if x=<> then <> else if x=<x ₁ ,x ₂ , ...,x _n > and 1≤n		
then <x _n ,x _{n-1} , ...,x ₁ > else ?		
rotl : x ==> if x=<> then <> else if x=<y> then <y> else if x=<x ₁ ,x ₂ , ...,x _n > and 1≤n		
then <x ₂ ,x ₃ , ...,x _n ,x ₁ > else ?		
rotr : x ==> if x=<> then <> else if x=<y> then <y> else if x=<x ₁ ,x ₂ , ...,x _n > and 1≤n		
then <x _n ,x ₁ ,x ₂ , ...,x _{n-2} ,x _{n-1} > else ?		
concat : x ==> if x=<<x ₁ ,x ₂ , ...,x _n >,<y ₁ ,y ₂ , ...,y _m >,...,<z ₁ ,z ₂ , ...,z _l >> and 1≤n,m,...,l		
then <x ₁ ,x ₂ , ...,x _n ,y ₁ ,y ₂ , ...,y _m , ...,z ₁ ,z ₂ , ...,z _l > else ?		
pair : x ==> if x=<x ₁ ,x ₂ , ...,x _n > and n is even then <<x ₁ ,x ₂ >,...,<x _{n-1} ,x _n >>		
else if x=<x ₁ ,x ₂ , ...,x _n > and n is odd then <<x ₁ ,x ₂ >,...,<x _n >> else ?		
split : x ==> if x=<y> then <<y>,>> else if x=<x ₁ ,x ₂ , ...,x _n > and 1≤n		
then <<x ₁ ,x ₂ , ...,x _{floor(n/2)} >,<x _{floor(n/2)+1} , ...,x _n >> else ?		

distl (distr) pairs from the left (the right) the first (second) argument with each element of the second (first) argument, apndl (apndr) appends the first (second) element to the end of the list given by the second (first) argument, trans treats the argument lists as the rows of a matrix and transposes this matrix, reverse reverses a list, rotl (rotr) does a left (right) circular shift on a list, concat concatenates two lists, pair pairs up the elements in a list, and split divides a list into nearly equal halves. Examples follow:

distl : <u <r m 2>>	==>	<<u r> <u m> <u 2>>
distr : <<i u voo> doo>	==>	<<i doo> <u doo> <voo doo>>
apndr : <<hi di> ho>	==>	<hi di ho>
apndl : <r <2 d 2>>	==>	<r 2 d 2>
trans : <<1 2 3> <a b c>>	==>	<<1 a> <2 b> <3 c>>
reverse : <to be or not to be>	==>	<be to not or be to>
rotl : <10 20 30 40>	==>	<20 30 40 10>
rotr : <<1 2> <3 4> <x y z>>	==>	<<x y z> <1 2> <3 4>>
concat : <<> <y z> <u q>>	==>	<x y z u q>
pair : <romeo juliet samson delila>	==>	<<romeo juliet> <samson delila>>
split : <hair1 hair2 h3 h4 h5>	==>	<<hair1 hair2 h3> <h4 h5>>

Predicates:

atom : x ==> if x is an atom then T else if x≠? then F else ?		
null : x ==> if x=<> then T else if x≠? then F else ?		
length : x ==> if x=<> then 0 else if x = <x ₁ ,x ₂ , ...,x _n > then n else ?		

atom tests whether the argument is an atom, null tests whether it is the empty list, and length returns the length of a list. Examples follow:

atom : firefox	==>	T
null : <not empty at all>	==>	F
length : <tell how long <am I>>	==>	4

Arithmetic and logical operations:

$+ : x \implies$ if $x = <a, b>$ and a and b are numbers then $a+b$ else ?
 (and similarly for $-$, $*$, $/$, $<$, $>$, \leq , \geq , $=$, \neq)
 $\text{xor} : x \implies$ if $x = <a, b>$ and a and b are boolean then $\text{xor}(a, b)$ else ?
 (and similarly for and, or, not, nand, and nor)

These operators compute the standard arithmetic and logical functions. Examples follow:

$+ : <2 3>$	\implies	5
$- : <4 13>$	\implies	-9
$* : <13 11>$	\implies	143
$/ : <81 3>$	\implies	27
$< : <100 999>$	\implies	T
$= : <k k>$	\implies	T

Functional forms, combinators, and conditional:

```

f1 @ f2 : x ==> f1 : (f2 : x)
[expr1,expr2, ...,exprn] : x ==> <expr1:x,expr2:x, ...,exprn:x>
& expr : x ==> if x = <> then <> else if x = <x1,x2, ...,xn> and 1 ≤ n
        then <expr : x1, expr : x2, ...,expr : xn> else ?
! expr : x ==> if x = <y> then <y> else if x = <x1,x2, ...,xn> and 2 ≤ n
        then expr : <x1, !expr:<x2, ...,xn>> else ?
\expr : x ==> if x = <y> then <y> else if x = <x1,x2, ...,xn> and 2 ≤ n
        then expr : <expr : <x1, ...,xn-1>,xn> else ?
| expr : x ==> if x = <y> then <y> else if x = <x1,x2, ...,xn> and 2 ≤ n
        then expr : <|expr : <x1, ...,xfloor(n/2)>, |expr : <xfloor(n/2)>>> else ?
(C -> A ; B) : x ==> if C : x = <> then A : x else if C : x = <> then B : x else ?

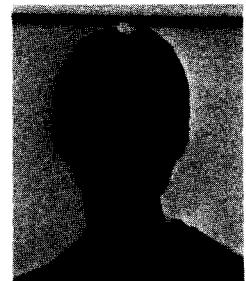
```

@ denotes functional composition and is by far the most common FP operator. Square brackets ([and]) denote list construction. The functor & maps the given function onto each element of the argument list. !, \, and | invoke several kinds of parallel tree evaluation, while the construct (C -> A;B) is the FP conditional operator: if C evaluates to true, A is evaluated and returned, otherwise B is evaluated and returned. Examples follow:

$* @ \text{tlr} : <3 2 1>$	\implies	6
$[+@[2,2],1,2] : <b 1>$	\implies	<2 b 1>
$\& \text{atom} : <\text{yes } <\text{no}> \text{ <<} \text{maybe} \text{>>>}$	\implies	<T F F>
$! * : <1 2 3 4 5 6>$	\implies	720
$ \text{id} : <1 2 3 4 5 6>$	\implies	<<<1 2> 3> <<4 5> 6>>
$(\text{atom} -> \%1 ; \%1) : \text{oxygen}$	\implies	1

ABOUT THE AUTHOR

Gabriel Robins was born in New York in February of 1962, and was raised in Israel. In 1974 he moved back to New York, and in 1976 he moved to Los Angeles, where he is currently residing. He received his B.S. in Mathematics and Computer Science from UCLA in 1983 and his M.S.E. in Computer Science from Princeton University in 1985. For the last several years he has been working part-time for Information Science Institute at Marina Del Rey, where he has contributed to the development of the knowledge representation language NIKL [1] and has most recently designed and implemented a portable software tool for the pictorial display of graphs [2,3,4], now commercially marketed by ExperTelligence Inc. Gabriel truly enjoys teaching, and is now a teaching assistant in the UCLA CS department, for the automata theory and formal languages course. In addition, he has developed with Dr. Eli Gafni a computational geometry course, given at UCLA for the first time in the Winter of 1988. Gabriel's interests include algorithms, data structures, complexity, graph theory, formal languages, computational geometry, user interfaces, and distributed & parallel algorithms. Presently he is working with Dr. Eli Gafni on some aspects of distributed protocols.



- [1] Kaczmarek, T., Bates, R., and Robins, G., "Recent Developments in NIKL", *Proceedings of the Fifth National Conference on Artificial Intelligence*, American Association of Artificial Intelligence, Philadelphia, August 1986.
- [2] Robins, G., "The ISI Grapher: a Portable Tool for Displaying Graphs Pictorially", invited talk in *Symbolikka '87*, Helsinki, Finland, August 17-18, 1987 (reprinted in *Multicomputer Vision*, Levialdi, S., Chapter 12, Academic Press, London, 1988).
- [3] Robins, G., "The ISI Grapher Manual", *ISI Technical Manual/Report ISI/TM-88-197*, USC/Information Sciences Institute, Marina Del Rey, February 1988.
- [4] Robins, G., "Applications of the ISI Grapher", *Proceedings of the Artificial Intelligence and Advanced Computer Technology Conference*, Long Beach, California, May 1988.

STATEMENT OF PRINCIPLES

The *Computer Science Department Quarterly* is issued three times a year by the Computer Science Department of the University of California at Los Angeles. It is directed to prospective students, current students, and alumni of the Computer Science Department as well as to interested members of the academic and scientific community. The purpose of the Quarterly is two-fold: (1) to present in a systematic form information regarding faculty resources, academic programs, research activities, and administrative matters relevant to the Computer Science Department; (2) to present short technical contributions, including tutorial articles and reports on recent research within the Computer Science Department.

In order to fulfill the first of the above objectives, each of the four annual issues emphasizes a different aspect of the Computer Science Department activities.

The Fall/Winter quarter issue is devoted to a description of the faculty resources of the Computer Science Department including Visiting Scholars, Teaching Assistants, Research Staff, and Fellowship Holders, with attention to special honors and awards received by faculty members, bibliographical citations, and lists of graduate degree recipients during the preceding calendar year.

The Spring quarter issue emphasizes academic programs, stressing particularly the requirements for the M.S. and Ph.D. degrees as well as descriptions of Computer Science Department course offerings and recommended course groupings.

The Summer quarter issue includes brief descriptions of the research interests and projects of the Computer Science faculty and reports on some of the sponsored and unsponsored research programs within the Computer Science Department.

**Computer Science Department
School of Engineering and Applied Science
3731 Boelter Hall
University of California
Los Angeles, CA 90024-1596**

**Non-Profit Org.
U.S. Postage
PAID
UCLA**

**Return Postage Guaranteed
Address Correction Requested**

