Adventures in the Forth Dimension

Julian V. Noble

Forth is a language that, for most programmers, is like the back side of the Moon: they know it is there but have never laid eyes on it. Yet Forth is mature (about as old as Lisp or C); ubiquitous in everyday applications such as automated tellers, package tracking, airline reservations, telephone switching systems, and microcontrollers everywhere; and built into millions of personal computers and workstations in the guise of IEEE Open Firmware. The page description language PostScript also has its roots in Forth.

Since Forth is, like the game of Go, simple yet deep, scientists and engineers should be better acquainted with it. This article provides an overview of the language with emphasis on availability[1], ease of use[2], structure, readability, and features considered important or essential in modern languages. I will also discuss execution speed and code size. To compensate for the omissions imposed by brevity I include an Appendix of Web links to sources of further information.

Forth is usually considered a language for embedded programming and control. However it is frequently applied in the less familiar contexts of robotics, data processing, security systems, music, and number crunching. Though not often taught in academic settings, Forth is a good pedagogical language since it offers access to all the computer's facilities and few predefined amenities. A student learns intimately how a heap or linked list works because he must construct it himself. I use Forth as the main illustration language in my course, *Computational Methods of Physics*[3], because pro-

---

[1] see Taygeta Scientific Inc. in the Appendix

[2] see Forth tutorials in the Appendix

[3] http://www.phys.virginia.edu/classes/551/

gram flow is clear, development is rapid, and lessons are reinforced through immediate feedback.

Forth came to my attention in the mid-1980's, at a time when my mainstay, Fortran, had let me down. Since it was easy to learn and a great time saver in program development, I came to prefer it to Fortran or C for virtually all computational work. For applications such as symbolic manipulation I have found it to be in a class with Lisp. One reason I chose Forth over C or Pascal is that it is easily extended to include both data structures and arithmetic operators for complex arithmetic (or quaternions, or anything else one might conceive).

The Forth language is simultaneously interpretive and compiled. Its grammar is simple: programs consist of sequences of words and numbers, separated by spaces (blank characters). The words are subroutines that are executed when named. Since Forth is interactive, a program can be typed in at the terminal; it can also be read from a file *via* commands like

```
FLOAD filename.ext   or   INCLUDE filename.ext
```

I find the root of the transcendental equation $xe^x - 1 = 0$ by typing in the phrases

```
: f1    FDUP  FEXP  F*  1.0e0  F-  ;  ok
USE( f1 0e0 2e0 1e-6 )FALSI   FS. 5.67143E-1  ok
```

The first line defines the function $f_1(x) = xe^x - 1$; the second line seeks the root of `f1` in the range $[0.0, 2.0]$ with precision 10^{-6} and displays it to the screen. The (previously compiled) subroutine `USE(` passes the function name, `)FALSI` finds the root by the hybrid *regula falsi* algorithm, and `FS.` displays the result. "ok" is what the interpreter says if all went well.

Forth is incrementally compiled: each subroutine—such as `f1` above—compiles as it is entered. Since most Forth subroutines can be executed interactively, we can test a new definition as soon as it has been entered, without writing a separate test program or elaborate scaffolding. This accelerates the program development cycle relative to traditional languages like C or Fortran. My own experience indicates a 10-fold speedup over Fortran;

2

Forth programmers who are also C-adept claim a similar ratio over development in C.

Two key themes underlie Forth's simplicity of structure and grammar. First, Forth simplifies communication between subroutines by providing direct access to the cpu stack. Conventional languages like C or Pascal communicate *via* temporary "stack frames" that hold arguments and parameters. A calling program constructs the frame, saves the current state of the system and the argument list on it, then transfers control to the subprogram. Upon exit the system state is restored and the frame deconstructed. This elaborate protocol can double or treble the execution time of a short subroutine. To reduce significantly the subroutine calling overhead, programs in such languages must reduce the number of subroutines called. That is, languages with significant communication overhead encourage programs featuring lengthy, versatile subroutines that perform several actions. Further speed optimizations are obtained by "in-lining" functions (rather than treating them as subroutines), "unrolling" loops, and other hallowed tricks that trade memory usage for execution speed.

Forth subroutines expect their arguments on the stack, and leave their results on the stack as well. This leads naturally to a "reverse-Polish" or "postfix" programming style, in which arguments precede functions, as with the HP family of RPN calculators. Using the stack directly to communicate between subprograms eliminates the stack-frame overhead and encourages programs consisting of small subroutines performing single functions. Program flow becomes clearer and debugging easier. Repeated operations are generally turned into subroutines and given their own names. This *factoring* process, the opposite of in-lining, shortens source codes and their corresponding executables.

The second key to simplicity is that Forth recognizes only numbers and subroutine names. That is, operators, compiler directives, control structures, data structures, commands, functions—in short, any programming constructs yet conceived—are subroutines that do their jobs when executed. For example, a named `VARIABLE` is a subroutine that places the address of the storage allocated to it on the stack. We define and test a new `VARIABLE` *via* the input lines
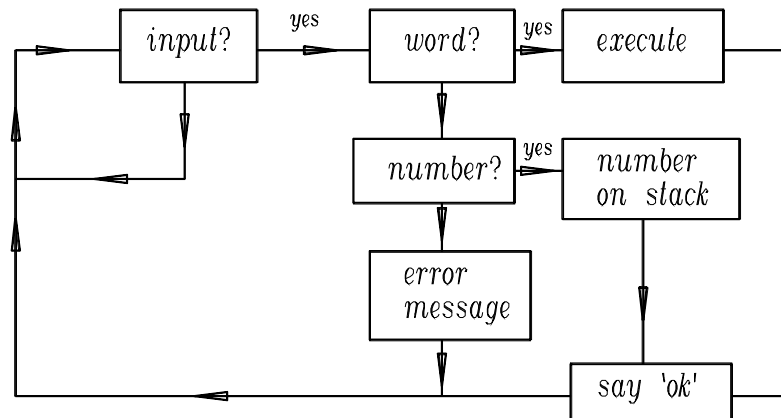
```
VARIABLE x  ok
7 x !  ok
x @  .  7  ok
```

The first line defines a variable named x. The second line inputs the string 7 and converts it to an integer (which is pushed on the stack), then executes x (which puts its own address on the stack). The subroutine ! (pronounced "store") expects a memory address on top of the stack and a number just below it. It stores the number to the address, consuming both stack items. The third line executes the new variable x (putting its address on the stack) and then executes the subroutine @ ("fetch"), which consumes an address on the stack and replaces it with the contents of that memory location. The subroutine "dot" (.) consumes the number on top of the stack—in this case, 7—and displays it on the standard output device—in this case the CRT. Similarly, the phrase

```
17 CONSTANT Seventeen  ok
Seventeen .  17  ok
```

defines (and tests) the subroutine Seventeen that, when executed, puts its contents (the integer 17) on the stack. This is what we want a CONSTANT to do.

A Forth interpreter is a simple endless loop that awaits input as shown in Fig. 1:

Each incoming item of text, delimited by spaces, is interpreted as follows:

1. Search the dictionary to see whether it is the name of a previously defined subroutine (a *word* in Forth parlance, since it is kept in a dictionary).

2. If found, pass its code address (or *execution token*) to the subroutine **EXECUTE** which executes the found subroutine, then issue the message "ok".

3. If not found, the subroutine **NUMBER** tries to convert the text to an integer (in the current base), and if successful pushes the result onto the stack and says "ok". (In a system extended to include floating point operations, **NUMBER** is modified so it recognizes a floating point number—such as 2e0—converts it appropriately, and pushes it onto the floating point stack. Of course this requires the system to be set for base-ten arithmetic, since 2e0 is a perfectly acceptable hexadecimal integer.)

4. If the text cannot be interpreted as a number, issue an error message.

5. Return to the beginning of the loop and await the next item of input.

Consider the input line

```
    3 4 5  *  +  .   23  ok
```

What happened? The interpreter interpreted sequentially the strings 3, 4 and
5 as integers and placed them on the stack (the most recent, 5, is on top).
The subroutine ∗ then multiplied the top two integers (4 and 5), consuming
them and leaving their product (20) on the stack. Next the subroutine +
added the top two integers (now 3 and 20), consuming them and leaving their
sum (23). And finally, the subroutine "dot" (.) consumed the top integer
on the stack and printed it to the display. Having executed this sequence of
operations successfully without running into undefined text, the interpreter
said "ok". That is, the interpreter acts like an RPN calculator.

In most compiled languages the compiler is a standalone program that in-
gests a file of source code, links it appropriately to external libraries, and out-
puts an executable file. The user controls some aspects of compilation with
"compiler directives" and/or "switches" (input with the source file or com-
mand line) but cannot extend or modify the compiler itself without rewrit-
ing it completely. In contrast, the Forth compiler is part of the interpreter.
There are actually two modes, *interpret* and *compile*, determined by the cur-
rent value of the system variable `STATE`. Figure 1 shows what happens with
`STATE` set to *interpret*. Setting it to *compile* modifies the behavior of the
routines `EXECUTE` and `NUMBER`: the former records the address of a routine
instead of executing it; whereas `NUMBER` encapsulates the number with code
that pushes it on the stack when the new subroutine runs.

Compilation in Forth uses a method called *threading*, that is most easily
explained with an example. Recall the interpretive example above, which
contained the sequence ∗ +. If this sequence appeared many times in a
program it might be useful to define a distinct subroutine that performs
both actions:

```
    : *+    *  + ;   ok
```

The subroutine "colon" (:) interprets the next word in the input stream
(in this case, ∗+) as the name of a new subroutine, creates a new dictionary
entry with that name, then switches to *compile* mode. `EXECUTE` records (that
is, *compiles*) pointers to ∗ and + in a list in the body of the new subroutine.

6

The final semicolon installs some terminating code, then switches the system back to *interpret* mode. The "ok" shows the process encountered no snags. We test interactively by entering

```
3 4 5  *+    .   23  ok
```

Since we get the correct answer and an "ok", `*+` is now a compiled, tested and debugged subroutine. This example illustrates several important points:

1. Forth does not restrict the characters forming the name of a subroutine. Good Forth style employs telegraphic names like `*+` that suggest what the subroutine does—one route to programming clarity.

2. A (compiled) Forth subroutine looks and acts just like interpreted Forth: a sequence of numeric literals and subroutine names that is executed left to right.

3. A new subroutine compiles as it is entered, therefore can be tested immediately.

4. Testing is simple: one places appropriate arguments on the stack, then invokes the subroutine. If it produces the right answer (without unexpected side effects) it can be considered debugged.

5. A Forth subroutine, once defined, becomes part of the language, on a par with the routines that came predefined with the system. There is no distinction between "user-defined" and "system"—which is why Forth is so easily (and often!) modified by its users.

6. A Forth program consists of subroutine definitions, ending with the subroutine that will actually perform the actions one wants. Thus a program to solve linear equations looks like

```
\ ... previous definitions of initialize,
\  triangularize, back_solve  and report

: }}solve   ( A[0]  V[0] --)  \ solve Ax = V; result in V
    initialize  triangularize  back_solve  report  ;
```

```
\ usage:  say  A{{ V{  }}solve
```

Conventional compilers are quite complex because they must recognize
control structures like `IF...ELSE..ENDIF` and decide what machine code to
generate. Forth control structures compile themselves: they are subroutines
that, like the terminating semicolon, always execute (rather than having their
addresses recorded when the system is in *compile* mode). The phrase

```
: +OR-    IF  +   ELSE  -  THEN    ; ok
```

defines a subroutine incorporating control structures. As we have seen, the
colon (:) initiates, and the semicolon (;) terminates, compilation of `+OR-`.
The routines `IF`, `ELSE` and `THEN` (the latter is what `ENDIF` is called in Forth)
lay down appropriate branching code. They use system pointers to com-
pute relative addresses to branch to, then insert these addresses back in the
branching code. (It is simpler than it sounds—these subroutines are defined
by only a line or two of Forth.) Testing, we have

```
3 5 TRUE  +OR- .  8  ok      \ added
3 5 FALSE +OR- .  -2  ok     \ subtracted
```

Control structures for indefinite (`BEGIN...UNTIL`, `BEGIN...WHILE...REPEAT`)
and definite (`DO...LOOP`) loops work more-or-less the same way.

Forth is a minimal language, which is why it has been implemented on so
many different kinds of computers. A Forth can be reduced to somewhere be-
tween 20 and 30 routines that must be defined in machine language; the rest
can be defined in terms of this basic kernel. The 1994 ANS Forth Standard[4]
lists 133 subroutines in the `CORE` wordset. This is the number required for a
system to advertise itself as ANS-compliant. The Standard specifies 372 sub-
routines altogether, mandating specific names and behaviors for single- and
double-length integer and floating point arithmetic; exponential, logarith-
mic, trigonometric and hyperbolic functions; file access; exception handling;

---

[4] ...available in various formats including HTML

generic keyboard and display I/O; generic memory management; and the built-in assembler, if provided (most Forths include one).

When I call Forth minimal, I mean it lacks predefined common constructs like linked lists or C-like *struct*s. However they can be added with little effort, since both commercial Forths and the vast body of public domain Forth programs (available on the Web) provide Standard implementations of arrays, queues, deques, stacks, heaps, lists, and so on. The Forth Scientific Library is a growing compendium of tested code for all the usual numerical algorithms. In other words, availability of code examples and libraries is really no hindrance to someone who wants to use Forth.

Forth provides features lacking in other languages, such as the ability to specify the current arithmetic base for number conversion. Some algorithms are more simply expressed in octal or hexadecimal arithmetic, so the ability to change bases freely is a boon. Standard Forth also permits executing strings of Forth code as though they were input from the command line or a file. This allows straightforward creation of macros, not to mention (safe) self-modifying code. (The latter is usually considered unsafe practice, but it can be very useful in artificial intelligence programming or language translation.)

Conspicuously absent from the 1994 ANS Standard are standardized names and behaviors for complex arithmetic, graphics, GUI construction, or port access. This is hardly surprising since cross-platform standards for such things do not exist for any language. (I do not wish to imply that no Forths provide such features—many do. It is just that they have not yet been standardized and are not portable.)

One of Forth's nicest features is that it hides nothing. Compiler words like : and `CODE` are comprised of components programmers can use freely to modify or extend the compiler; or for that matter, to define multiple compilers for specialized tasks. Using the dictionary look-up subroutine "tick" (') and the subroutine "comma" (,) that stores an integer on the stack into the next unused memory cell we can, for example, construct jump tables for rapid execution of multiple-choice programs by finding and storing the execution tokens of the subroutines comprising the table[5].

---

[5]Two implementations of this idea can be found in *A Beginner's Guide to Forth* at the

Defining a specialized mini-compiler can be a useful and powerful technique. Examples from my own work are:

1. Self-executing tables of random variates from arbitrary distributions—useful in Monte-Carlo simulations [2].

2. A compact recursive-descent FORmula TRANslator (about 500 lines of code and comments; 727 including white-space lines, documentation and conditional compilation sections) to translate formulas to RPN Forth.

3. A terse compiler (22 lines) of two-dimensional state-transition tables into the finite state machines they represent [3]. This is such a clear, efficient and expressive way to program FSMs that it has been applied to text processing and language translation, to computer game programming, and to gas/oil pipeline controllers, to name only applications I know about. The FSM mini-compiler simplifies and clarifies programs with lots of decisions (such as the Boyer-Moore text-search algorithm, symbolic γ-matrix reduction, or multi-variate function minimization by the simplex algorithm).

I am often asked how Forth stacks up against other languages. This depends on the criteria of comparison. Forth is versatile; and easy to obtain, use and understand. To compare it with other languages requires me to address the areas of execution speed, development time, code size, structure, readability, and features considered important or essential in modern languages.

Forth codes typically run 3 to 10 times slower than (optimized) C or Fortran equivalents. In other words, unoptimized Forths, such as the public domain systems I have been using[6], are about as fast as unoptimized Lisp or C++. There are, however, several ways to achieve fast execution in Forth. Simplest is to invest in an optimizing native code compiler; there are several available whose execution of numeric-intensive software compares favorably

---

website http://www.phys.virginia.edu/classes/551/

[6]For the past several years I have used Gforth and Win32Forth. The former runs on Unix/Linux platforms as well as under MS Windows[TM]. Both are ANS compliant.

with good C compilers—that is, about 1.5 to 2× slower than hand-tuned machine code.

For such specialized applications as linear equations, differential equations or Fourier transforms, commercial packages like MatLab^{TM} provide optimized libraries of code subroutines. These libraries can be linked to and called from Forth programs running under Linux or Windows^{TM}, using standard methods [4]. To use external libraries one must of course know their matrix labelling and subroutine calling conventions.

Traditionally we optimize Forth for speed by identifying bottleneck subroutines using algorithmic analysis or a profiler, then rewriting the bottlenecks in assembler. Since Forths usually include assemblers, this is an easier route to take than it would be with other languages. In fact writing and testing assembly language subroutines is much easier in the Forth environment than in any other I know of [5] since neither test program nor linking step is required. A Forth subroutine defined in `CODE` is used the same way as its high level equivalent. For example, the innermost loop of the LU algorithm for linear equations (repeated twice in the Fortran subroutine `LUDCMP` [6]) is a good candidate to be factored out and defined as a separate `CODE` subroutine. The instructions it contains are the only ones executed $\mathcal{O}(N^3)$ times; this loop therefore dominates the asymptotic running time. These instructions are so simple (two fetches, a multiplication and a subtraction) it is easy to hand-code. Further optimization affects only those instructions executed N^2 times or fewer. Hence with one `CODE` definition the Forth version of LU solves large dense systems at the maximum speed attainable by a non-vectorized machine.

Most often, however, my goals are development speed and program correctness rather than execution speed *per se*. By these criteria Forth beats the other languages I know. A typical instance was a program I wrote to simulate vehicular rollover accidents using realistic forces. It took me less than a day to write and debug the code, and a couple of hours to add a graphical display. From prior experience, I do not think I could have achieved this with Fortran in less than a week. Similarly, having never written any sort of parser or compiler before, I created a usable recursive-descent FORmula TRANslator in less than one week.

Since I am a physicist who programs, not a professional programmer, I cannot speak for the latter group. However, professionals report similar experiences. For example, the IEEE Open Firmware Specification evolved from Mitch Bradley's (secret) use of Forth to develop hardware drivers at Sun Computer Corp. According to Bradley, development went much faster if he prototyped in Forth and converted to the in-house language afterward. Again, the (public domain) Windows™-compliant Win32Forth was written by Tom Zimmer because he found the Microsoft Windows™ SDK *cum* C++ too slow to meet a promised deadline. He delivered his code on time by first creating an object-oriented Forth, then using it to construct Windows™ programs.

Forth programs tend to be more compact than their equivalents in other languages. This is true of both source and executables. For example, the 16-bit DOS-based Forth (that was my mainstay before Windows™) had a 32 Kbyte executable. The Windows™ compatible Win32Forth has a 52 Kb executable and a 400 Kb runtime library, which sounds like a lot until one compares it with other Windows™-based PC applications (see table below):

Program	Size (Kb)	DLL's (Kb)
HSFORTH 4.2 (16-bit DOS)	32	
GForth 0.40 (32-bit DOS)	176	
Aztec Forth (32-bit DOS)	133	
Win32Forth 3.5	452	288
WinView	410	288
Turbo C++ Lite	834	
Ventura Publisher 4.2	997	1017
WordPerfect 6.1	3755	3029

When it comes to source code size, the source for WinView, an extensive Windows™ editor included with Win32Forth, is 268 Kb. The source for the assembler (for Intel cpu's and fpu's) is about 80 Kb. And the source for my FORmula TRANslator is 24 Kb (including extensive comments).

What about readability and maintainance? At one time structured programming was a central goal of computer pedagogy. Nicholas Wirth invented Pascal in reaction to "spaghetti code" produced by students. Wirth aimed to eliminate line labels and direct jumps (GOTOs), thereby forcing control

flow to be clear and direct and making spaghetti code impossible. Paradoxically, Forth is the only truly structured language left in common use today, although that was not its raison d'être. It contains neither GOTOs nor line labels. A Forth subroutine has a single entry and a single exit point, and (usually) performs a single job.

Like every language, Forth can be written obscurely. I have certainly seen plenty of underdocumented, badly formatted, badly factored code with poorly named and excessively verbose subroutines. Of course this is true of C, Fortran or any other language—it is as easy to produce "write-only" code as to write muddy prose. Standard Forth offers the usual remedies: comments, stack diagrams, sensible and telegraphic naming conventions, proper layout. Choosing memorable and descriptive names for subroutines (remember, *everything* in Forth is a subroutine, so this includes variables and constants too) can help a lot, as in the linear equations example. The names initialize, triangularize, backsolve, report portray the algorithm so clearly that comments add nothing. However Forth also makes available some unusual remedies, that with care can produce exceptionally clear programs. A self-executing jump-table (compiled by the mini compiler jtab: ... ;jtab) *looks* like a table:

```
5 long jtab:  CandyMachine
        | Snickers   \ item 0
        | Payday     \ item 1
        | M&Ms       \ item 2
        | Hershey    \ item 3
        | AlmondJoy  \ item 4
;jtab
```

—hardly any further comment or explanation is needed.

One aspect of Forth disconcerting for the newcomer is its lack of safety features. Some Forths perform rudimentary stack checks during compilation, but this is by no means mandated by the Standard and is provided at the discretion of the vendor. Since arrays in Forth are defined as needed by the programmer, they can incorporate bounds checking or not, as desired. The jump tables compiled by jtab: ... ;jtab prevent out-of-bounds indexing by clipping, as good a way as any.

The omission of array bounds checking is not mere hacker machismo, however. Bugs that cause memory leaks are much rarer in Forth than in C. Forth errors tend to involve the stack—removing too many items or too few. Obviously a word that uses the stack incorrectly inside a moderately long loop can crash with the greatest of ease. A simple discipline of keeping word definitions short enough to understand, commenting them thoroughly—especially their stack effects—and testing each word as it is defined, eliminates the bulk of these errors. That is, I do not consider Forth programs inherently unsafe, despite the dearth of safety nets. And one professional programmer I know specializes in safety-critical applications using Forth [7].

A propos of safety, what about debugging? "Serious" languages nowadays come with an "environment" that includes integrated editing and code debugging of various levels of sophistication. Although many Forths provide single-stepping code debuggers, and some provide integrated editors, experienced Forth programmers mostly do without these appurtenances. Again this has little to do with machismo, but a great deal to do with the way programs get written. Most subroutines are short enough to be correct the first time. They can be tested as they are entered, and quickly reveal errors, omissions or oversights. I have needed debuggers only a few times, usually because I let a subroutine get too verbose. Some languages permit "assertions" [8] that specify conditions—usually loop invariants—that must be satisfied at various points in the program. This is considered a key to guaranteeing program correctness. Assertions have been implemented in Standard Forth, but I think more for the challenge than because they were needed. They may be valuable in a large C program, but in Forth assertions seem like overkill.

Some consider modularity essential in "serious" languages. Languages that support modularity enable multi-programmer teams to develop different parts of a large code in separate modules, joining them only at the end. This programming paradigm is doubtless the source of the enormous improvements we have witnessed of late, in the quality, ease of use, and reliability of commercial applications. ;-)

Well, Forth can be made modular also. One of the most successful commercial Forths for large 16-bit DOS applications [9, 10] was designed around this concept. In fact, modularity—in whatever strength you need—is relatively easy to implement in Forth. The weakest form, suitable for most

14

applications, is based on ANS Forth's support for partitioning the dictionary into distinct *wordlists*. Repeated subroutine names ("name-space collisions") are not a problem because the compilation mechanism can be told which wordlists not to search. Many Forth programmers habitually distinguish public and private resources by defining the words `public` and `private`, that execute appropriate search-order switches.

Discussions in the `comp.lang.*` newsgroups often revolve around memory management, particularly "garbage collection". As most readers will know, the need for this arises from dynamic memory allocation within a heap. It is not obvious how best to de-allocate memory that is no longer needed, since it can be physically located amidst memory that must be retained; worse, it is not always easy to determine whether a given chunk is, in fact, ready for reclamation. Some languages—such as Lisp—are more afflicted by garbage collection problems than others. In Forth, I am happy to report, garbage collection is mainly a non-issue. We tend to eschew variables, leaving temporary items on the stack; this space is reclaimed immediately when a word exits. Memory allocation is directly under the programmer's control, so de-allocating it can be handled in the way best suited to the specific application.

Because Forth's programming philosophy runs to simple subroutines that do single tasks, arithmetic operators are not overloaded. The integer multiplication `*` is not the same as the floating point operator `F*`. Moreover since there is no telling what might be on the stack at a given moment, we do not expect `*` to recognize when a floating point number must be multiplied by an integer. To simplify some of my programs I once implemented a form of operator overloading to allow mixed real/complex arithmetic. Since the decisions take place at run time rather than compile time (late, not early binding) this slows program execution about 15%.

Forth may be the most portable language in use today, with at least one version available for almost every current platform (and most obsolete ones as well). I have yet to hear of a program written entirely in ANS Forth that fails to perform correctly on another platform running ANS-compliant Forth. A mostly compliant program using a few non-Standard or environmentally-dependent subroutines generally can be ported successfully with little effort.

Finally, there is the matter of the programming "style" that a given language encourages and/or supports. Raw Forth can look almost like assembly language. In fact Forth has often been dismissed as "nothing more than a high-level assembly language". But Forth's enormous extensibility has allowed it to take on many guises. Lisps, Prologs, BASICs, SmallTalk's and C's have been written in Forth, sometimes for serious reasons and sometimes just for fun. Standard Forth provides the tools to construct any linguistic paradigm one wishes, even object orientation with polymorphism and inheritance: numerous modern Forths include object-oriented extensions (but there is as yet no concensus in the Forth community either that such refinements are necessary or desirable; or, if they are, what the standard interface should be).

Many of my own programs look like the evaluation of arithmetic formulae, *a la* Fortran, because I use a FORmula TRANslator (the subroutine `f"` reads a string up to a closing `"` and translates it to Forth on the fly). For example, here is how I translate a short Fortran subroutine [11] to Forth:

```fortran
      SUBROUTINE TRIDAG(A,B,C,R,U,N)
      PARAMETER (NMAX=100)
      DIMENSION GAM(NMAX),A(N),B(N),
#            C(N),R(N),U(N)
      IF(B(1).EQ.0.)PAUSE
      BET=B(1)
      U(1)=R(1)/BET
      DO 11 J=2,N
        GAM(J)=C(J-1)/BET
        BET=B(J)-A(J)*GAM(J)
        IF(BET.EQ.0.)PAUSE
        U(J)=(R(J)-A(J)*U(J-1))/BET
11    CONTINUE
      DO 12 J=N-1,1,-1
        U(J)=U(J)-GAM(J+1)*U(J+1)
12    CONTINUE
      RETURN
      END
```

```forth
include arrays.f
include ftran111.f

100 VALUE Nmax

Nmax long 1 FLOATS 1array a{     \ input array
Nmax long 1 FLOATS 1array b{     \ as 3 vectors
Nmax long 1 FLOATS 1array c{

0 VALUE aa{   0 VALUE bb{   0 VALUE cc{   0 VALUE NN

Nmax long 1 FLOATS 1array r{     \ inhomogeneous term

Nmax long 1 FLOATS 1array L{     \ diagonal
Nmax long 1 FLOATS 1array U{     \ lower subdiagonal
Nmax long 1 FLOATS 1array x{     \ solution vector

: }triangularize    ( a{ b{ c{ n --)
    TO NN    TO cc{   TO bb{   TO aa{
    f" bb{0}"
    FDUP  F0=  ABORT" Reduce # of equations by 1"
    f" L{0}"  F!
    f" U{0} = cc{0} / L{0}"
    NN 1-  0  DO  f" U{I} = cc{I} / L{I}"
                  f" L{I 1+} = bb{I 1+} - aa{I 1+} * U{I}"
    LOOP   ;

: }backsolve     ( r{ x{ n --)
    TO NN    TO aa{   TO bb{
    f" bb{0} = bb{0} / L{0}"
    NN 1 DO   f" bb{I} = (bb{I} - a{I}*bb{I 1-}) / L{I}"
    LOOP
    f" aa{NN 1-} = bb{NN 1-}"
    0 NN 2 -  DO   f" aa{I} = bb{I} - U{I}*aa{I 1+}"
    -1 +LOOP   ;

\ say:   a{ b{ c{ n }triangularize  r{ x{ n }backsolve
```

The Fortran looks terser but this is illusory: it lacks comments and white space, not to mention a program to use it. The Forth version runs as is.

The urge to use Forth as if it were something else afflicts programmers who have begun to appreciate Forth's power, but have not yet abandoned the habits of their previous language(s). I once attempted to write a full Fortran-to-Forth translator. However, as I soon discovered, good Fortran translated literally becomes terrible Forth, rather like Mark Twain's literal re-translation into English, of the French version of "The Celebrated Jumping Frog" [12]. Having learned my lesson, whenever I am tempted nowadays to change Forth's innate style, I emulate Bennet Cerf and lie down until the urge passes.

17

Appendix: Useful Web sites in the world of Forth

MPE Ltd. (http://www.mpeltd.demon.co.uk/)

FORTH, Inc. (http://www.forth.com/)

Online articles and books of interest:

C.H.Moore and G.C. Leach, *FORTH—A Language for Interactive Computing* (http://www.ultratechnology.com/f70c2.html)

Philip J. Koopman, Jr., *Stack Computers: the new wave* (http://www.cs.cmu.edu/~koopman/stack_computers/index.html)

Journal of Forth Application and Research (peer-reviewed online journal) (http://dec.bournemouth.ac.uk/forth/index.html)

The FORTH Research Page. Maintained and validated by Dr. Peter Knaggs (pjk@bcs.org.uk) (http://dec.bournemouth.ac.uk/forth/jfar/index.html)

Phil Burk, *Forth Tutorial* (http://www.softsynth.com/pforth/pf_tut.htm)

Taygeta Scientific Incorporated: large online collection of Forth info (http://www.taygeta.com/)

Forth Interest Group Home Page (http://www.forth.org/)

Bill Muench, *eForth: A simple model Forth system* (http://members.aol.com/forth/)

J.V. Noble, *Computational Methods of Physics* (http://www.phys.virginia.edu/classes/551/)
The programs alluded to in this article can be found there.

A Forth in Java (http://world.std.com/∼wware/agj/fj.html)

References

[1] J.V. Noble, "Avoid decisions" *Computers in Physics* **5** #4 (1991) 386.

[2] J.V. Noble, *Scientific Forth: a modern language for scientific computing* (Mechum Banks Publishing, Ivy, VA, 1992) pp. 61-64.

[3] J.V. Noble, "Finite state machines in Forth", *Journal of Forth Application and Research* **7** http://www.jfar.org/article001.html

[4] http://www.er.ele.tue.nl/emv/frames/EMV_hendrix.htm

[5] J.V. Noble, "A call to assembly" *Forth Dimensions* in press.

[6] W.H. Press, B.P. Flannery, S.A. Teukolsky and W.T. Vettering, *Numerical Recipes* (Cambridge University Press, New York, 1986) p. 35ff.

[7] http://www.amleth.demon.co.uk/

[8] Jon Bentley, *Programming pearls* (Addison-Wesley Publishing Co., Reading, MA, 1989) p. 42.

[9] Microprocessor Engineering, Ltd. (http://www.mpeltd.demon.co.uk/)

[10] Paul Frenger, "Learning Forth with Modular Forth", *ACM/SIGPLAN Notices* **35** #3 (2000) 25.

[11] W.H. Press, *et al.*, *ibid.*, p. 40.

[12] Samuel Clemens, *The family Mark Twain* (Dorset Press, New York, 1988), pp. 1163ff.

Julian Noble is Professor of Physics at the
Department of Physics
University of Virginia
P.O. Box 400714
Charlottesville, VA 22904-4714

He may be contacted at jvn@virginia.edu .

His interests are eclectic, both in and out of physics.
His teaching philosophy is "no black boxes".

Figure captions

1. Structure of the Forth interpreter