

The Evolution of Forth

Elizabeth D. Rather
 PORTH, Inc.
 111 N. Sepulveda Blvd., #300
 Manhattan Beach, CA 90256

Donald R. Colburn
 Creative Solutions, Inc.
 4701 Randolph Rd., Suite #12
 Rockville, MD 20853

Charles H. Moore
 Computer Cowboys
 410 Star Hill Road
 Woodside, CA 94062

Abstract

Forth is unique among programming languages in that its development and proliferation has been a grass-roots effort unsupported by any major corporate or academic sponsors. Originally conceived and developed by a single individual, its later development progressed under two significant influences: professional programmers who developed tools to solve application problems and then commercialized them, and the interests of hobbyists concerned with free distribution of Forth. These influences have produced a language markedly different from traditional programming languages.

Table of Contents

1	Chuck Moore's Programming Language.....	2
1.1	Early Development.....	2
1.2	Philosophy and Goals.....	4
2	Development and Dissemination.....	6
2.1	Forth at NIRA.....	7
2.2	Commercial Minitcomputer Systems.....	8
2.3	Early Microprocessor Systems.....	12
2.4	Language Definition.....	14
3	Forth Without Chuck Moore.....	22
3.1	The Forth Interest Group.....	22
3.2	Commercial and Public Domain Systems for Personal Computers.....	24
3.3	Embedded Systems.....	29
4	Hardware Implementations of Forth.....	31
5	Present and Future Directions.....	32
5.1	Standardization Efforts.....	33
5.2	Implementation Strategies.....	35
5.3	Object-oriented Extensions.....	36
6	A Posteriori Evaluation.....	37
6.1	Meeting Objectives.....	37
6.2	Major Contributions of Forth.....	38
6.3	Mistakes or Desired Changes.....	39
6.4	Problems.....	41
6.5	Implications for Current and Future Languages.....	41
	References.....	42
	Bibliography.....	45

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

HOPH-11/4/93/IMA,USA
 © 1993 ACM 0-89791-571-2/93/0004/0177...\$1.50

1. Chuck Moore's Programming Language

Forth was invented by Charles H. (Chuck) Moore. A direct outgrowth of Moore's work in the 1960's, the first program to be called Forth was written in about 1970 [Moore, 1970a]. This section covers the early work leading to Forth.

1.1 Early Development

Moore's programming career began in the late 1950's at the Smithsonian Astrophysical Observatory with programs to compute ephemerides, orbital elements, satellite station positions, etc. [Moore, 1958], [Veis, 1960]. His source code filled two card trays. To minimize recompiling this large program, he developed a simple interpreter to read cards controlling the program. This enabled him to compose different equations for several satellites without recompiling. This interpreter featured several commands and concepts that survived into modern Forth, principally a command to read "words" separated by spaces and one to convert numbers from external to internal form, plus an **IF ... ELSE** construct. He found free-form input to be both more efficient (smaller and faster code) and reliable than the more common Fortran practice of formatting into specific columns, which had resulted in numerous re-runs due to mis-aligned columns.

In 1961, Moore received his BA in Physics from MIT and entered graduate school at Stanford. He also took a part-time programming position at the Stanford Linear Accelerator (SLAC), writing code to optimize beam steering for the (then) pending two-mile electron accelerator, using an extension of some of his prior work with least-squares fitting. A key outgrowth of this work was a program called **CURVE**, coded in Algol (1964), a general-purpose non-linear differential-corrections data fitting program. To control this program, he used an enhanced version of his interpreter, extended to manage a push-down stack for parameter passing, variables (with the ability to explicitly fetch and store values), arithmetic and comparison operators, and the ability to define and interpret procedures.

In 1965, he moved to New York City to become a free-lance programmer. Working in Fortran, Algol, Jovial, PL/I and various assemblers, he continued to use his interpreter as much as possible, literally carrying around his card deck and recoding it as necessary. Minicomputers appeared in the late 60's, and with them teletype terminals, for which Moore added operators to manage character input and output. One project involved writing a Fortran-Algol translator and file-editing utilities. This reinforced for him the value of spaces between words, which were not required in Fortran source.

Newly married and seeking a small town environment, Moore joined Mohasco Industries in Amsterdam, NY, in 1968. Here he developed computer graphics programs for an IBM 1130 minicomputer with a 2250

graphic display. This computer had a 16-bit CPU, 8K RAM, his first disk, keyboard, printer, card reader/punch (used as disk backup!), and Fortran compiler. He added a cross-assembler to his program to generate code for the 2250, as well as a primitive editor and source-management tools. This system could draw animated 3-D images, at a time when IBM's software for that configuration drew only static 2-D images. For fun, he also wrote a version of Spacwar, an early video game, and converted his Algol Chess program into the new language, now (for the first time) called FORTH. He was impressed by how much simpler it became.

The name FORTH was intended to suggest software for the fourth (next) generation computers, which Moore saw as being characterized by distributed small computers. The operating system he used at the time restricted file names to five characters, so the "rj" was discarded. FORTH was spelled in upper case until the late 70's because of the prevalence of upper-case-only I/O devices. The usage "Forth" was generally adopted when lower case became widely available, because the word was not an acronym.

Moore found the Forth-based 1130 environment for programming the 2250 superior to the Fortran environment in which the 1130 software was developed, so he extended it into an 1130 compiler. This added looping commands, the concept of keeping source in 1024-byte blocks and tools for managing them, and most of the compiler features we recognize in Forth today.

Most important, there was now a dictionary. Procedures now had names, and the interpreter searched a linked list of names for a match. Names were compiled with a count and three characters, a practice learned from the compiler writers of Stanford and which prevailed in Forth until the 1980's. Within a dictionary entry was a "code field" containing the address of code to be executed for that routine. This was an indirect threaded code implementation (see Section 5.2) and was in use five years before Dewar's paper on indirect threaded coded appeared in *Communications of the ACM* [Dewar 1975]. The use of indirect threaded code was an important innovation, since an indirect jump was the only overhead once a word had been found. Dictionary entries could consist either of pointers to other "high level" routines or of machine instructions.

Finally, in order to provide a simple mechanism for nesting routines, a second stack called the "return stack" was added. The benefit of having a stack reserved for return addresses was that the other stack could be used freely for parameter passing, without having to be "balanced" before and after calls.

The first paper on Forth was written at Mohasco [Moore, 1970a].

In 1970 Mohasco assigned Moore to an ambitious project involving a new Univac 1108 handling a network of leased lines for an order-entry system. He

ported Forth onto the 1108, and arranged for it to interface to COBOL modules that did the transaction processing. The 1108 Forth was coded in assembler. It buffered input and output messages and shared the CPU among tasks handling each line. It also interpreted the input and executed the appropriate COBOL modules. This version of Forth added mechanisms for defining and managing tasks, and also an efficient scheme for managing disk block buffers similar to schemes in use today.

Unfortunately, an economic downturn led Mohasco to cancel the 1108 project before completion. Moore immediately gave notice, then wrote an angry poem and a book on Forth [Moore, 1970b] that was never published. It described how to develop Forth software and encouraged simplicity and innovation.

1.2 Philosophy and Goals

To Moore, Forth was a personal response to his frustration with existing software tools, which he viewed as a sort of "tower of Babel" [Moore, 1970a]:

The software provided with large computers supplies a hierarchy of languages: the assembler defines the language for describing the compiler and supervisor; the supervisor the language for job control; the compiler the language for application programs; the application program the language for its input. The user may not know, or know of, all these languages: but they are there. They stand between him and his computer, imposing their restrictions on what he can do and what it will cost.

And cost it does, for this vast hierarchy of languages requires a huge investment of man and machine time to produce, and an equally large effort to maintain. The cost of documenting these programs and of reading the documentation is enormous. And after all this effort the programs are still full of bugs, awkward to use and satisfying to no one.

Moore conceived of Forth as replacing the entire "vast hierarchy" with a single layer, requiring only two elements: a programmer-to-Forth interface, consisting of minimal documentation (minimal because the interface should be simple and natural), and the Forth-machine interface, consisting of the program itself.

His view was entirely personal, considering his own needs in the light of his own experience. The following excerpts from his unpublished book [Moore, 1970b], describe this view:

I've written many programs over the years. I've tried to write *good* programs, and I've observed the manner in which I write them rather critically. My goal has been to decrease the effort required and increase the quality produced.

In the course of these observations, I've found myself making the same mistakes repeatedly. Mistakes that are obvious in retrospect, but difficult to recognize in context. I thought that if I wrote a prescription for programming, I could at least remind myself of problems. And if the result is of value to me, it should be of value to others....

Above all, his guiding principle, which he called the "Basic Principle," was, "Keep it simple!" Throughout his career he has observed this principle with religious dedication.

As the number of capabilities you add to a program increases, the complexity of the program increases exponentially. The problem of maintaining compatibility among these capabilities, to say nothing of some sort of internal consistency in the program, can easily get out of hand. You can avoid this if you apply the Basic Principle. You may be acquainted with an operating system that ignored the Basic Principle.

It is very hard to apply. All the pressures, internal and external, conspire to add features to your program. After all, it only takes a half-dozen instructions, so why not? The only opposing pressure is the Basic Principle, and if you ignore it, there is no opposing pressure.

The main enemy of simplicity was, in his view, the siren call of generality that led programmers to attempt to speculate on future needs and provide for them. So he added a corollary to the Basic Principle: "Do not speculate!"

Do not put code in your program that *might* be used. Do not leave hooks on which you can hang extensions. The things you might want to do are infinite; that means that each has 0 probability of realization. If you need an extension later, you can code it later — and probably do a better job than if you did it now. And if someone else adds the extension, will he notice the hooks you left? Will you document this aspect of your program?

This approach flew in the face of accepted practice then as now. A second corollary was even more heretical: "Do it yourself!"

The conventional approach, enforced to a greater or lesser extent, is that you shall use a standard subroutine. I say that you should write your own subroutines.

Before you can write your own subroutines, you have to know how. This means, to be practical, that you have written it before; which makes it difficult to get started. But give it a try. After writing the same subroutine a dozen times on as many computers and languages, you'll be pretty good at it.

Moore followed this to an astounding extent. Throughout the 70's, as he implemented Forth on 18 different CPUs (Table 1), he invariably wrote for each his own assembler, his own disk and terminal drivers, even his own

multiply and divide subroutines (on machines that required them, as many did). When there were manufacturer-supplied routines for these functions, he read them for ideas, but never used them verbatim. By knowing exactly how Forth would use these resources, by omitting hooks and generalities, and by sheer skill and experience (he speculated that most multiply/divide subroutines were written by someone who had never done one before and never would again), his versions were invariably smaller and faster, usually significantly so.

Moreover, he was never satisfied with his own solutions to problems. Revisiting a computer or an application after a few years, he often re-wrote key code routines. He never re-used his own code without re-examining it for possible improvements. This later became a source of frustration to Rather, who, as the marketing arm of FORTH, Inc. (see Section 2.2), often bid jobs on the assumption that since Moore had just done a similar project this one would be easy — only to watch helplessly as he tore up all his past code and started over.

Today, Moore is designing Forth-based microprocessors using his own Forth-based CAD system, which he has re-written (and sometimes rebuilt, with his own hardware) almost continuously since 1979.

Moore considered himself primarily an applications programmer, and regarded this as a high calling. He perceived that "systems programmers" who built tools for "applications programmers" to use had a patronizing attitude toward their constituents. He felt that he had spent a great fraction of his professional life trying to work around barriers erected by systems programmers to protect the system from programmers and programmers from themselves, and he resolved that Forth would be different. Forth was designed for a programmer who was intelligent, highly skilled and professional; it was intended to empower, not constrain.

The net result of Moore's philosophy was a system that was small, simple, clean — and extremely flexible: in order to put this philosophy into practice, flexible software is essential. The reason people leave hooks for future extensions is that it's generally too difficult and time-consuming to re-implement something when requirements change. Moore saw a clear distinction between being able to teach a computer to do "anything" (using simple, flexible tools) and attempting to enable it to do "everything" with a huge, general-purpose OS. Committing himself to the former, he provided himself with the ideal toolset to follow his vision.

2. Development and Dissemination

By the early 1970's, Forth had reached a level of maturity that not only enabled it to be used in significant applications, but that attracted the attention of other programmers and organizations. Responding to their needs, Moore

implemented it on more computers and adapted it to handle ever larger classes of applications.

2.1 Forth at NRAO

Moore developed the first complete, stand-alone implementation of Forth in 1971 for the 11-meter radio telescope operated by the National Radio Astronomy Observatory (NRAO) at Kitt Peak, Arizona. This system ran on two early minicomputers (a 16 KB DDP-116 and a 32 KB H316) joined by a serial link. Both a multiprogrammed system and a multiprocessor system (in that both computers shared responsibility for controlling the telescope and its scientific instruments), it was responsible for pointing and tracking the telescope, collecting data and recording it on magnetic tape, and supporting an interactive graphics terminal on which an astronomer could analyze previously recorded data. The multiprogrammed nature of the system allowed all these functions to be performed concurrently, without timing conflicts or other interference.

The system was also unique for that time in that all software development took place on the minis themselves, using magnetic tape for source. Not only did these Forth systems support application development, they even supported themselves. Forth itself was written in Forth, using a "metacompiler" to generate a new system kernel when needed.

To place these software capabilities in context, it's important to realize that manufacturer-supplied system software for these early minicomputers was extremely primitive. The main tools were cross-assemblers and FORTRAN cross-compilers running on mainframes (although the FORTRAN cross-compilers were too inefficient to do anything complex, given the tiny memories on the target machines). On-line programming support was limited to assemblers loaded from paper tape, with source maintained on paper tape. Digital Equipment Corp. had just announced its RT-11 OS for its PDP-11 line, which offered limited foreground-background operation; no form of concurrency was available for the H316 family. Multiuser operation of the sort that enabled NRAO's astronomers to graphically analyze data while an operator controlled the telescope and live data was flowing in, was unheard-of.

Edward K. Conklin, head of the Tucson division of NRAO which operated the 11-meter telescope, found it difficult to maintain the software since Moore was based at NRAO's headquarters in Charlottesville, VA. So in 1971 he brought in Elizabeth Rather, a systems analyst at the University of Arizona, to provide local support on a part time basis. Rather was appalled to find this critical system written in a unique language, undocumented and known to only one human. Her instinctive reaction was to re-write the whole thing in FORTRAN to get it under control. Alas, however, there was neither time nor budget for this, so she set out to learn and document the system as best she could.

After about two months, Rather began to realize that something extraordinary was happening: despite the incredibly primitive nature of the on-line computers, despite the weirdness of the language, despite the lack of any local experts or resources, she could accomplish more in the few hours she spent on the Forth computers once a week than the entire rest of the week when she had virtually unlimited access to several large mainframes.

She wondered why. The obvious answer seemed to lie in the interactive nature of Forth. The programmer's attention is never broken by the procedural overhead of opening and closing files, loading and running compilers, linkers, loaders, debuggers, etc. But there's more to it than that. For example, all the tools used by Forth's OS, compiler, and other internal functions are available to the programmer. And, as Chuck Moore intended, its constraints are minimal and its attitude is permissive. Forth devotees still love to debate the source and magnitude of such productivity increases!

Rather immediately left the University and began working for NRAO jointly with Kitt Peak National Observatory (KPNO), an optical observatory with which NRAO shared facilities, maintaining the Forth system for NRAO and developing one for KPNO (which was later used on KPNO's 156" Mayall telescope and other instruments [Phys. Sci. 1975]). During the next two years she wrote the first Forth manual [Rather, 1972] and gave a number of papers and colloquia within the observatory and related astronomical organizations [Moore, 1974a].

In 1973 Moore and Rather replaced the twin-computer system by a single disk-based PDP-11 computer [Moore, 1974a&b]. This was a multi-user system, supporting four terminals in addition to the tasks controlling the telescope and taking data. It was so successful that the control portions of it were still in use in 1991 (data acquisition and analysis functions are more dependent on experimental equipment and techniques, which have changed radically over the years). The system was so advanced that astronomers from all over the world began asking for copies of the software. Versions were installed at Steward Observatory, MIT, Imperial College (London), the Cerro Tololo (Chile) Inter-American Observatory, and the University of Utrecht (Netherlands). Its use spread rapidly, and in 1976 Forth was adopted as a standard language by the International Astronomical Union.

2.2 Commercial Minicomputer Systems

Following completion of the upgraded system in 1973, Moore and his colleagues Rather and Conklin formed FORTH, Inc. to explore commercial uses of the language. FORTH, Inc. developed multiuser versions of Forth [Rather, 1976a] for most of the minicomputers then in use (see Table 1), selling these as components of custom applications in a widely diverse market, ranging from data base applications to scientific applications such as image processing. The minicomputers and applications of 70's provided the

available, and manufacturer-supplied executives such as DEC's RT-11 supported foreground-background operation. Multiuser systems were also being common: a PDP-11 or Nova could be expected to support up to eight users, although the performance in a system with eight active users was poor.

On this hardware, Moore's Forth systems offered an integrated development toolkit including interactive access to an assembler, editor and the high-level Forth language, combined with a multitasking, multiuser operating environment supporting 64 users without visible degradation, all resident without run-time overlays.

Although time-critical portions of the system were written in assembler, as most applications required very high performance, Moore could port an entire Forth development environment to a new computer in about two weeks. He achieved this by writing Forth in Forth — any Forth computer could generate Forth for another, given the target system's assembler and code for about 60 primitives. Since the first step in a port was designing and writing the target assembler, it is possible that Moore has written more assemblers for different processors than anyone else.

Being able to port the system easily to new architectures was important, since the minicomputer market was extremely fragmented. A large number of CPUs was available, and each was supported by a large number of possible disk controller and drive combinations. Today, by contrast, the microcomputer market is dominated by a very short list of processor families, and adherence to *de-facto* standards such as the PC/AT is the norm.

Installations were done on site, since it was impractical to ship the minicomputers. When LSI-11's first became available, Moore bought one and mounted it in a carry-on suitcase, with a single 8" floppy drive in a second suitcase. This portable personal computer accompanied him everywhere until 1982, acting as a "friendly" host for generating new Forths.

2.2.2 Application requirements

If the principal environmental constraints were memory limitations and a need to serve a broad spectrum of CPU architectures, the application requirements were dominated by a need for performance. Here are some of the principal application areas in which Forth achieved success in this period:

1. Commercial/business data base systems: First developed for Cybek Corp. under the guidance of Arthur A. Gravina, these systems supported multiple terminals on a Data General Nova, handling high-speed transaction processing. The first was written for Vernon Graphics, Inc., a service bureau to Pacific Telephone, in 1974. It supported 32 terminals processing transactions against a 300 MB data base. In its first week the system handled over 100,000 transactions a day

environment in which Forth developed and stabilized, to the extent that all of the innovations contributed by independent implementors in the years that followed represented relatively minor variants on this theme. Because of this, we shall take a close look at the design and structure of these systems.

Year	Model	Customer	Forth Applications
1970-71	Honeywell H316	National Radio Astronomy Observatory (NRAO)	Data acquisition, on-line analysis w/ graphics terminal
1971	Honeywell DDP116	NRAO	Radio telescope control
1971-2	IBM 370/30	NRAO	Data analysis
1972	Varian 620	Kitt Peak National Observatory (KPNO)	Optical telescope control and instrumentation
1972	HP2100	KPNO	Instrumentation
1972-3	Modcomp	NRAO	Data analysis
1973	PDP-11	NRAO	Radio telescope control, data acquisition, analysis, graphics
1973	DGNova	Steward Observatory	Data acquisition and analysis
1974	SPC-16	Steward Observatory	Ground control of balloon-borne telescope
1975	SDS920	Aerospace Corp.	Antenna control
1975	Prime	Gen1 Dynamics, Pomona	Environmental controls
1976	Four-Phase	Source Data Systems	Data entry and data base management
1977	Interdata Series 32	County of Alameda, CA	Data base management
1977	CALSI-4	MICOA	Business systems
1978	Honeywell Level 6	Source Data Systems	Data entry and data base management
1978	Intel 8086	Aydin Controls	Graphics and Image Processing
1980	Raytheon 100	PTS- American Airlines	Airline display and workstations

Table 1

Table showing computers for which Chuck Moore personally implemented Forth systems. In 1978, his implementations of Forth on the Level 6 and 8086 represented the first resident software on both CPUs, anticipating their manufacturers' systems by many months.

2.2.1 Environmental constraints

Minicomputers of the 1970's were much less powerful than the smallest microcomputers of today. In the first half of the decade not all systems even had disks — 1/2" tape was often the only mass storage available. Memory sizes ranged from 16 to 64 Kbytes, although the latter were considered large. In the early 70's, most programming for minis was done in assembly language. By the middle of the decade compilers for Fortran and BASIC were

These applications are characterized by high data rates, as much as 20 KHz in some cases, which really strained the CPU speed of the processors available. Fast interrupt response was essential, along with high-speed multitasking to allow data acquisition to proceed concurrently with operator activity and instrument control.

2.2.3 Influences

The evolution of Forth prior to 1978 was completely dominated by Moore himself. As we have seen, Moore was and is a fanatic minimalist, dedicated to the principle of zero-based design in which every feature and every instruction must justify its existence or be ruthlessly scrapped.

Moore originally developed the system for his own use. It surprised him a little to find that Rather and the other early users also liked it and found it enhanced their productivity as much as it did his. But even after the formation of FORTH, Inc. and its open marketing of the system, the selection and design of support tools and the general programming interface was dominated by his personal tastes.

Moore was working primarily as a consultant, supported by others within FORTH, Inc., installing a Forth system on a customer's computer as the first step in developing a custom application. Since the customer was primarily interested in the application, it was imperative that the port be completed quickly and inexpensively. The extreme simplicity of Forth made this possible without compromising the performance of the application.

Each of these projects contributed its own lessons, tools and techniques. Moore carried microfiche listings of all previous projects in his briefcase, and often referred to them to get the code for some unique primitive or driver from the past. Frequently used words might become a standard fixture of the system. Also, improved techniques for solving common problems were integrated into the system.

This pattern of continual evolution created customer support headaches for FORTH, Inc., however, as no two installed systems were the same. In most cases the installation included a five-day Forth programming course taught by Rather, who had to check every evening to make sure that the system still behaved the way it was being taught.

2.3 Early Microprocessor Systems

In 1976, Robert O. Winder, of RCA's Semiconductor Division engaged FORTH, Inc. to implement Forth on its new CDP-1802 8-bit microprocessor [Rather, 1976b], [Electronics, 1976]. The new product, called "microFORTH," was subsequently implemented on the Intel 8080, Motorola 6800 and Zilog Z80, and sold by FORTH, Inc. as an off-the-shelf product. microFORTH was

(40,000 was the requirement). The system was subsequently upgraded to support 64 terminals and a 600 MB data base, with no discernable degradation in response times, which remained under one second.

Cybek subsequently marketed this system for business applications in banking and hospital management; its current version is marketed by a division of McDonnell Douglas. A similar effort by Source Data Systems in Iowa produced a multi-terminal data-entry system marketed by NCR Corp. for hospital management and similar applications.

The performance of such a system is overwhelmingly dominated by operating system issues, principally the ability of the native Forth block-based file system to read and write data files very quickly.

2. Image Processing: FORTH, Inc. developed a series of image processing applications for the Naval Weapons Research Center, NASA's Goddard Space Flight Center, the Royal Greenwich Observatory in England, and others. Central to these was a need for performing standardized operations (e.g., enhancement, windowing, etc.) on images residing on different kinds of hardware. The approach taken included many features now associated with object oriented programming: encapsulation (the basic object was an "image," with characteristic parameters and methods), inheritance (you could add new images that would inherit characteristics of previously defined classes of images) and dynamic binding of manipulation methods. Moore, the principal architect of this approach, was unaware of any academic work in this area. Striving to achieve the same goals as later OOPS writers, he independently derived similar solutions.

Image processing systems are also distinguished by a need to manipulate and move large quantities of data very fast: a 512x512x16 image, for example, occupies 512 KB. In addition to the high-speed disk performance that characterized Forth data base systems, these also required fast processing speed and the ability to handle algorithms such as FFTs. As many minicomputers lacked hardware floating-point arithmetic, Forth included flexible integer and fixed-point fraction operators, as well as specialized array primitives.

3. Instrumentation and control: Forth was first developed and used for this purpose at NRAO, and Forth is widely used for instrumentation and controls today. FORTH, Inc. did several more astronomical systems (for the Universities of Wyoming [Gehrz, 1978], Minnesota, Hawaii and Illinois; Cal Tech; plus the Royal Greenwich Observatory and St Andrews University in the UK). In addition, a number of commercial instrument manufacturers such as Princeton Applied Research (now a division of EG&G) and Nicolet Instruments adopted Forth as a language for internal development.

successfully used in numerous embedded microprocessor instrumentation and control applications in the United States, Britain and Japan.

2.3.1 Environment and applications

microFORTH was FORTH, Inc.'s first experience with off-the-shelf, mail-order software packages; the minicomputer systems were all installed on-site. The mail-order operation was made possible by the rapid standardization of the industry on 8" "IBM-format" floppy disks, and the relatively small number of development systems for each CPU type.

These microprocessors were all 8-bit devices, typically with 16K bytes of memory in the development system. The target systems were usually custom boards (although Intel's Single Board Computer series quickly became popular) and the software was expected to run from PROM in an embedded environment without disk or (usually) terminal. This was significantly different from the minicomputer environment, where there was always a disk, and a program was expected to run on the same (or identical) computer as the one used for development.

Most microprocessor manufacturers offered development platforms consisting of the same microprocessor as in the target, up to 64K bytes of RAM, a serial line for a terminal, a parallel printer port, and two 8" floppy disk drives. Software support was mainly assembler, although Intel soon introduced PL/M. In-circuit emulators and separate utilities were introduced for debugging.

microFORTH was principally marketed as an interactive alternative to assembler which, unlike PL/M, was available across most microprocessor families and therefore offered a higher degree of transportability.

2.3.2 Language definition

Following some initial experimentation with 8-bit stack width and 128-byte block buffers, it was quickly decided to maintain the same basic internal architecture as on the minicomputer systems. The organization of the program changed significantly, however.

microFORTH came with a target nucleus designed to run from PROM. This nucleus was only 1K in size, containing primitives such as single-precision arithmetic and other very basic functions. The development environment supported writing and testing code interactively, and then compiling a version of that code designed to mate to the run-time nucleus. A version of **VARIABLE** was provided to support segregated ROM/RAM data space (**CONSTANTS** were in PROM), and defining words were adapted so that user-defined structures could be made to reside in either. And whereas previously **VARIABLEs** could be initialized at compile time, that capability was removed,

as it's difficult to initialize target RAM when a ROM is being compiled without setting up a "shadow" table; ROM space was considered too precious for that.

The multiprogramming support was initially stripped out, although it later came back using a new, faster task-swapping algorithm, and the data base tools vanished completely.

FORTH, Inc. never released the metacompiler used to generate Forth on new minicomputer CPUs. A variant of this metacompiler became an integral part of microFORTH, however, as it was used to generate the ROMable code for the target application. This was significant, as we shall see in the next section.

2.3.3 Influences

The principal architect of microFORTH was Dean Sanderson. Although Sanderson worked closely with Moore and shared most of his basic philosophies, differences in style and approach were inevitable. But the major new influence came from the broader customer base that resulted from the wider marketing of microFORTH. It was customer pressure that brought back multiprogramming, and this larger customer base also caused standards groups to form.

2.4 Language Definition

The commercial mini and microcomputer implementations produced by FORTH, Inc. in the early- and mid-1970's for the first time encapsulated the principles and elements of Forth as it is used today. For this reason, we shall summarize these briefly.

2.4.1 Design principles

Much as algebra was the "metaphor" for FORTRAN, Forth was conceived on the model of English prose (though some have suggested that its postfix notation tends to resemble verb-at-the-end languages such as German). Its elements ("words") are named data items (roughly equivalent to nouns), named procedures (equivalent to verbs), and defining words (special kinds of verbs capable of creating data items with customized characteristics). Words may be defined in terms of previously defined words or in machine code (using the embedded assembler).

Forth "words" are functionally analogous to subroutines in other languages. They are also equivalent to commands in other languages — Forth blurs the distinction between linguistic elements and functional elements.

Words are referenced (either from the keyboard or in program source) by name. As a result, the term "word" is applied both to program (and

linguistic) units and to their text names. In parsing text, Forth considers a word to be any string of characters bounded by spaces (or "white space" characters in some file-based systems). Except for these, there are no special characters that cannot be included in a word or start a word, although many programming teams adopt naming conventions to improve readability. Words encountered in text fall into three categories: defined words (i.e., Forth routines), numbers, and undefined words.

There are no explicit typing mechanisms in Forth, a feature which sometimes surprises newcomers, but is generally admired by experienced Forth programmers.

2.4.2. Structured programming disciplines

Architecturally, Forth words adhere strictly to the principles of "structured programming" as articulated by Dijkstra (e.g., Dijkstra, 1970) and "modular programming" [Parnas, 1972]. These principles may be summarized as follows:

- Every program is described as a linear sequence of self-contained modules.
- A module has one entry point and one exit point and ideally performs one function, given a set of inputs and a set of outputs.
- A module can contain:
 - references to other modules
 - decision structures (**IF THEN** statements)
 - looping structures

Top-down design and bottom-up coding and testing are strongly encouraged by Forth's structure.

As was the case with Moore's independent development of OOPs-like features in his image processing system, Moore was unfamiliar with the contemporary literature on structured programming. These principles were first called to his attention in 1973 by Rather, who received several comments on the apparent relationship between Forth and structured programming in seminars she was giving on Forth. On reading one of Dijkstra's papers, Moore observed, "it just seems like good programming practice to me."

In fact, advanced Forth programmers with knowledge of the underlying implementation know ways of "cheating," but such practices are frowned upon, and definitely not supported or encouraged by the structure of the language.

2.4.3. Elements of Forth

Moore's Forth systems of the early 1970's were built on a nucleus of only 4K bytes. This tiny program included disk (or tape) and terminal drivers and the

ability to search and build the dictionary. This nucleus was then used to compile from source the balance of the programming environment, including the assembler, editor, multiuser support and several hundred general commands. Booting the system, including compiling most of it from source into executable form, took only a few seconds.

A metacompiler, also written in Forth, was used to compile the nucleus. The entire source for the system was about 40 pages long.

These systems were "native," that is, running without any host OS or executive. This was a necessity in the early days, as OSs weren't available. Later, it was regarded as a significant advantage, as I/O services in a native Forth environment were much faster than could be supplied by a general purpose OS.

The principal elements of Forth will be discussed briefly in the sections that follow.

2.2.4.1 Dictionary

A Forth program is organized into an extensible dictionary that occupies almost all the memory used by the system. The dictionary is classically implemented as a linked list of variable-length items, each of which defines a word. The content of each definition depends upon the type of word (data item, constant, sequence of operations, etc.). On multi-user Forth systems individual users may have private dictionaries, each of which is connected to a shared, re-entrant system dictionary.

2.2.4.2 Push-down stacks

Forth maintains two push-down stacks, or LIFO lists (on a multiprogrammed version, a pair for each task). These are used to pass data between Forth words and for controlling logical flow. A stack contains one-cell items, where a cell is 16 bits wide on 8-bit and 16-bit computers and 32 bits wide on most implementations for 32-bit processors such as the 680x0 family. Extended-precision numbers occupy two stack positions, with the most significant part on top. Items on either stack may be addresses or data items of various kinds. Stacks are of indefinite size, and usually grow towards low memory.

Forth's explicit use of stacks leads to a "postfix" notation in which operands precede operators. Since results of operations are left on the stack, operations may be strung together effortlessly, and there is little need to define variables to use for temporary storage.

2.2.4.3 Interpreters

Forth is an interpretive system, in that program execution is typically controlled by a small machine-code routine (often only two or three instructions)

interpreting lists of pointers or tokens for abstract machine functions. This architecture is much faster than classical interpreters, as used in BASIC and PROLOG for example, enabling it to perform satisfactorily in the real-time applications for which it was designed.

This internal engine is often referred to as the "inner" or "address" interpreter, as distinct from Forth's more traditional text interpreter which processes source and user input. The text interpreter extracts strings separated by spaces from the terminal or mass storage, looking each word up in the dictionary. If a word is found it is executed by invoking the address interpreter, which processes a string of addresses compiled in a word definition by executing the definition pointed to by each. The text is not stored in memory, even in condensed form. If a word is not found, the system attempts to convert it as a number and push it onto the stack. If number conversion fails (due to a non-numeric character), the interpreter aborts with an error message.

The address interpreter has two important properties. First, it is fast, often requiring as few as one or two machine instructions per address. Second, it makes Forth definitions extremely compact, as each reference requires only one cell (or computer word; Forth users prefer to avoid the use of "word" as a hardware unit because of its use to denote an element in the language). In contrast, a subroutine call constructed by most compilers requires instructions for handling the calling sequence before and after a `CALL` or `JSR` instruction and address, and typically save and restore registers within the subroutine. Forth's stack architecture obviates the need for an explicit calling sequence, and most implementations make global register assignments in which certain system state variables are assigned to dedicated registers, and all other registers are designated scratch registers for use in code words.

2.2.4.4 Assembler

Most Forth systems include a macro assembler for the CPU on which they run. When using `CODE` the programmer has full control over the CPU, as with any other assembler, and `CODE` definitions run at full machine speed. The assembler lets the programmer use explicit CPU-dependent code in manageable pieces with machine-independent interfacing conventions. To move an application to a different processor requires re-coding only the `CODE` words, which interact with other Forth words in exactly the same manner.

Forth assemblers feature an unusual design, which has two goals: (1) to improve transportability between processors by standardizing assembler notation as much as possible without impairing the programmer's control of the processor, and (2) to yield a compact assembler that can be resident at all times to facilitate interactive programming and debugging.

In a classical Forth assembler, the `op-code` itself is a Forth word which assembles the instruction according to operands passed on the stack giving

the addressing information. This leads to a format in which the addressing mode specifiers precede the `op-code` (consistent with the postfix notation used elsewhere in Forth). Moore also standardized notation for addressing modes, although he usually used the manufacturer's instruction mnemonics. Registers were generally referred to by number, except for registers assigned to key internal system functions. For example, the stack pointer is usually in a register called `s`. One would address the second item on a two-byte wide stack using the phrase `2 s`.

Forth assemblers support structured programming in the same way that high-level Forth does. Arbitrary branching to labelled locations is discouraged; on the other hand, structures such as `BEGIN ... UNTIL` and `IF ... ELSE ... THEN` are available in the assembler (implemented as macros that assemble appropriate conditional and unconditional branches). Such structures are easy to implement because the stack is available during assembly to carry addressing information.

Conventional assemblers leave the code in a file, which must be integrated with code in files from high-level language compilers (if any) by a linker before the resultant program can be loaded into memory for testing. The resident Forth assembler assembles the code directly into memory in executable form, thus avoiding the linking step.

The Forth assembler is used to write short, named routines that function just like high-level Forth words: when the name of the routine is invoked, it will be executed. Like other Forth routines, code routines expect their arguments on the stack and leave their results there. Within code a programmer may refer to constants (to get a value), variables (to get an address) or other defined data types. Code routines may be called from high-level definitions just as other Forth words, but do not themselves call high-level or code definitions.

These features enable Forth programmers to write code in short, easily testable modules that are automatically integrated into an application. Programming is fully structured, with consistent rules of usage and user interface for both assembler and high-level programming. Words are tested incrementally, while the desired behavior is fresh in the programmer's mind. Most new words can be tested simply by placing input values on the stack, typing the word to be tested and validating the result left on the stack by displaying it.

The result is complete control of the computer, high performance where needed, and overall shortening of development time due to interactive programming at all levels.

2.2.4.5 Disk support

Classical Forth divides mass storage into "blocks" of 1024 bytes each. The block size was chosen as a convenient standard across disks whose sector sizes vary. At least two block buffers are maintained in memory, and the block

management algorithm makes it appear that all blocks are in memory at all times. The command `n BLOCK` returns the memory address of block `n`, having read it if necessary. A buffer whose contents is changed is marked so that when it needs to be reused its block is automatically written out. This algorithm provides a convenient form of virtual memory for data and source storage, with a minimum number of physical disk accesses required. `FORTH, Inc.`'s data base applications build data files out of blocks, with a file defined as spanning a specified range of blocks; data access is through operations performed against named fields within selected files.

In native Forths, the block system is both fast and reliable, as the disk driver computes the physical address of the block from its number — no directory is required. In disk-intensive applications, performance can be enhanced by adding more buffers, so more blocks will be found in memory; the buffers become a disk cache.

In the 1980's, Forth systems became available running under conventional OSs, as we shall see. Many of these support blocks within host OS files, although some have abandoned blocks altogether. As blocks provide a compatible means of accessing mass storage across both native and non-native systems, `ANS Forth` (Section 5.1) requires that blocks be available if any mass storage support is available.

2.2.4.6 Multiprogramming

The earliest Forth systems supported multiprogramming, in that the computer could execute multiple concurrent program sequences. In 1973, Moore extended this capability to support multiple users, each with a terminal and independent sub-dictionaries and stacks. The entity executing one of these program sequences or supporting a user is referred to as a task. Many of today's Forths support multiprogramming, and most of these use variants of Moore's approach.

This approach allocates CPU time using a cooperative, non-preemptive algorithm: a task relinquishes the CPU while awaiting completion of an I/O operation or upon use of the word `PAUSE`, which relinquishes the CPU for exactly one lap around the round-robin task queue.

Moore's systems used interrupts for I/O. Interrupts were directly vectored to the response code using an assembler macro, without intervention by the Forth executive. Interrupt code performed only the most time-critical operations (e.g., read a number, increment a counter), then reenabled the task that had been suspended pending the interrupt. The task would actually resume operation the next time it was encountered in the round-robin task loop, at which time it would complete any high-level processing occasioned by the event and continue its work.

In theory this non-preemptive algorithm is vulnerable to a task monopolizing the CPU with logically or computationally intensive activity, but in practice real-time systems are so dominated by I/O that this is rarely a problem. Where CPU-intensive operations do occur, `PAUSE` is used to "tune" performance.

Consultant Bill Cox pointed out [Cox, 1987] that a non-preemptive algorithm such as this has several advantages. First, the task scheduler itself is simpler and faster, taking as little as one machine instruction per task. Second, since a task is suspended only at known, well-defined times, it has less "context" to be saved and restored, so the context-switch itself is faster. Third, task code can be written with the knowledge of exactly when the task does or does not control the CPU, and management of shared resources is considerably simplified. Cox compared the performance of several real-time OSs; the results are given in Table 2.

Event:	VRTX	OS9	PDOS	POLYFORTH
Interrupt response	91	43.75	93.4	7.0
Context switch	128	186.25	93.4	36
Suspend task	180	316.25	184.7	6.8
Copy memory (80 bytes)		212.5		97

Table 2

Performance comparisons of several real-time OSs on a M68010 [Cox, 1987]. Times are averages, given in μ s. Times were normalized to a 10 MHz 68010. `POLYFORTH`'s use of non-preemptive task scheduling accounts for its performance advantage.

Tasks were constructed when the system was booted, and each was given a fixed memory allocation adequate to the functions it was intended to perform. As re-booting took only a few seconds, it was easy to reconfigure a task.

2.2.4.7 Computation

Until the late 1970's, few minicomputers offered floating point arithmetic — indeed, many lacked hardware multiply and divide. From the beginning, however, Forth was used for computationally intensive work. Controlling the radio telescope, for example, required converting wanted positions from the celestial coordinates in which astronomical objects are located to an azimuth/elevation coordinate system once per second and interpolating intermediate positions five times per second, with data acquisition and operator activity proceeding concurrently.

Moore's approach was to build into Forth the ability to manipulate integers effectively. For example, the command `* /` multiplies two single-cell integers and divides by a third, with a double-length intermediate product. This reflects the way most multiply and divide machine instructions work, and enables calculations such as:

12345 355 113 */

This phrase multiplies 12345 by the ratio 355/113, which represents π with an error of 8.5×10^{-6} [Brodie, 1981]. The ability to multiply by a ratio is ideal for calibration and scaling, as well as rational approximations. Similarly, the word /MOD performs a single division, returning both the quotient and remainder. A rich set of single, double and mixed-precision operations such as these make integer arithmetic much more usable than it is in most languages.

Moore expressed angles internally as 14-bit, 15-bit or 30-bit fixed-point binary fractions. He provided a set of primitives to convert to and from angle formats (e.g., dd:mm:ss), and a math library supporting transcendental functions for these formats based largely on algorithms from Hart [1968]. Operations such as the Fast Fourier Transform were provided in some applications, built on specialized primitives supporting complex numbers as scaled integer pairs.

Today fast floating-point processors are common. Many Forths support floating point, as does ANS Forth. But in many cases, such as embedded systems on simple microcontrollers, Forth's integer arithmetic still provides simpler, faster solutions.

2.2.4.8 Data types

Perhaps nowhere was Moore's personal philosophy more in evidence than in his approach to data typing. Basically, he wanted to assume full responsibility for manipulating data objects in whatever way he wished. If pressed on this point, he would say, "If I want to add 1 to the letter A, it's none of the compiler's business to tell me I can't."

Standard Forth words support single and double-precision **CONSTANTS**, which return their values on the stack, and **VARIABLES**, which return a pointer. **CREATE** names the beginning of a data region in which space can be reserved. The pointer returned by a **CREATED** entity can be incremented to index into an array. The nature of the values kept in constants and variables was entirely arbitrary; there is normally no explicit type checking. Strings are normally kept in memory with their length in the first byte. The address of this structure, or the address and length of the actual string, can be passed on the stack.

CONSTANT, **VARIABLE** and **CREATE** are "defining words" that is, they define new words with characteristic behaviors. Forth also provides tools to enable the programmer to build new defining words, specifying a custom behavior both at compile time (e.g., setting up and initializing a table) and run-time (e.g., accepting an index and automatically applying it to the base address of the structure).

3. Forth Without Chuck Moore

microFORTH was heavily marketed, and attracted a lot of attention in the late '70s. One side effect of this was the growth of an active and enthusiastic group of hobbyists who fell in love with Forth. In their wake came new companies marketing versions of Forth in competition with FORTH, Inc. At the same time, Moore himself was becoming increasingly drawn toward hardware implementations of Forth, and less involved in software production at FORTH, Inc. (which he left in 1982 to pursue his hardware interests full time). In this section we examine the development of Forth under these diverse new influences.

3.1 The Forth Interest Group

In the late 1970's, Northern California was afire with the early rumblings of the Computer Revolution. Groups of interested individuals such as the "Home Brew Computer Club" were meeting to share interests and experiences. Magazines such as *Radio Electronics* published step-by-step instructions on how to build your own video display terminal, and even how to build your own microcomputer system.

Due to the high cost of memory and low level of VLSI integration, typical "homebrew" computers were very resource-constrained environments. Echoing back to the first generation computers, there was insufficient memory to concurrently support an editor, assembler and linker. Mass storage was slow and expensive, so many homebrew systems used paper tape or audio cassette tapes for I/O. Although some BASIC language products were available, they were typically very slow, and incapable of supporting significant programs. The stage was thus set for something else to meet the expanding needs of these hardy explorers and "early adopters."

Forth had been born and bred to exploit the minimal facilities of resource-constrained systems. It carried neither the excess baggage of a general solution nor a requirement for an existing file or operating system or significant mass storage. As Forth was used to tackle more and more difficult embedded computer applications, it started to claim the attention of the Northern California homebrew computer enthusiasts.

Bill Ragsdale, a successful Bay Area security system manufacturer, became aware of the benefits of microFORTH, and in 1978 asked FORTH, Inc. to produce a version of microFORTH for the 6502. FORTH, Inc. declined, seeing much less market demand for microFORTH on the 6502 than the more popular 8080, Z80 and 6800 CPUs.

Ragsdale then looked for someone with the knowledge of microFORTH and intimate familiarity with the 6502 to port a version of microFORTH to the 6502. He found Maj. Robert Selzer, who had used microFORTH for an AMI

6800 development system on an Army project and was privately developing a standalone editor/assembler/linker package for the 6502. Selzer wrote a 6502 Forth assembler, and used the Army's microFORTH metacompiler to target compile the first 6502 stand-alone Forth for the Jolt single board computer.

Selzer and Ragsdale subsequently made substantial modifications and improvements to the model, including exploitation of page zero and stack-implicit addressing architectural features in the 6502. Many of the enhancements that characterized the later public-domain versions were made during this period, including variable-length name fields and modifications to the dictionary linked-list threading. A metacompiler on the Jolt could target a significantly changed kernel to a higher address in memory. A replacement bootable image would then be recompiled by the new kernel into the lower boot address, which could then be written out to disk. At this point, Ragsdale had a system with which to meet his professional needs for embedded security systems.

During this period the Forth Interest Group (FIG) was started by Ragsdale, Kim Harris, John James, David Boulton, Dave Bengel, Tom Olsen and Dave Wyland [FIG 1978]. They introduced the concept of a "FIG Forth Model," a publicly available Forth system that could be implemented on popular computer architectures.

The FIG Forth Model was derived from Ragsdale's 6502 system. In order to simplify publication and rapid implementation across a wide variety of architectures, a translator was written to convert Forth metacompiler source code into text that, when input to a standard 6502 assembler, would replicate the original kernel image. In this way, neither the metacompiler nor its source code needed to be published. This is an important point. Forth metacompilation is a difficult process to understand completely. It requires the direct manipulation of three distinct execution phases and object areas, and is not something that a casual user wanted or needed.

By publishing assembler listings, the Forth Interest Group was able to encapsulate a Forth run-time environment in a manner that could be easily replicated and/or translated to the assembly language of a different computer architecture. It was the intention of the original team of implementors to thus stimulate the development of compatible Forth systems and the appearance of new vendors of Forth products.

After the 6502 FIG Model was published, FIG implementors published compatible versions for the 8080 and 6800 microcomputers and the PDP-11 and Computer Automation minicomputers. Over the years, volunteers added other platforms and documentation. The 1982 *Forth Encyclopedia* by Mitch Derick and Linda Baker [Derick, 1982] provided an exhaustive 333-page manual on FIG Forth, with flow charts of most words. In 1983 an ad in *Forth Dimensions*, the FIG newsletter [FIG, 1983], listed: RCA 1802, 8080, FACE,

6502, 8086/88, 6800, 6809, 9900, Nova, Eclipse, VAX, Alpha Micro, Apple II, 68000, PDF11/LSI11 and Z80.

Today there are several thousand members of the Forth Interest Group in over fifteen countries. Since 1980, FIG has sponsored an annual conference called FORML (Forth Modification Laboratory), an educational forum for sharing and discussing new or unproven proposals intended to benefit Forth, and for discussion of technical aspects of Forth. Its proceedings are available from the Forth Interest Group (P. O. Box 8231, San Jose, CA 95155).

3.2 Commercial and Public Domain Systems for Personal Computers

Apple Computer grew out of the bubbling computer enthusiasm in the San Francisco Bay area, and with it a whole new generation of resource-constrained computers. Although BASIC was available in ROM, Forth was used to write a number of popular text editors and games on the Apple II, allowing resident development of significant programs within its scarce memory and disk constraints. It is hard now, with ubiquitous megabytes of memory and disk, to imagine what it was like to develop significant programs on a 40-column wide screen within 16K of memory and 100K of disk storage.

Vendors of low cost Forth systems sprang up almost overnight, each supporting their favorite personal computer, most of them basing their systems on the FIG model. In 1979, for example, Miller Microcomputer Services announced MMSFORTH for the TRS-80 [TRS-80, 1979], and by 1980 *Computerworld* reported [Taylor, 1980] that MMS had over 100 user groups for its product.

When IBM entered the personal computer business with their original PC product offering, they chose to distribute a version of the popular Apple II text editor EasyWriter, written in Forth, as an IBM product. Laboratory Microsystems (LMI) introduced a commercial IBM PC Forth system in 1982. Numerous commercial and public domain Forth products followed, and significant software product development began.

Following its introduction of the first commercial Forth for the IBM-PC, LMI has maintained a continuing strategy of producing cutting-edge Forth systems for the PC, including a 32-bit real-mode implementation (February 1983), an OS/2-based Forth (February 1988) and a Windows version (1992). Along the way LMI's founder, Ray Duncan, became an acknowledged authority on Microsoft OSs (e.g. Dugan, 1988).

FORTH, Inc.'s PC offering was polyFORTH, which combined the multiuser support and data base tools of its minicomputer products with the ROMable architecture of microFORTH. By 1984 FORTH, Inc. was supporting up to 16 users on a PC with no visible degradation, running polyFORTH first as a native OS and later as a co-resident OS with MS-DOS. By the late 80's

polyFORTH users such as NCR were supporting as many as 150 users on a single 80386-based PC.

System(s)	Company	Primary Products & Markets
CFORTH-83, Fortinacs, SunForth	Bradley Forthware	Portable Forth written in C; versions for Atari, Macintosh, Sun; consulting and services related to the Sun Microsystems Open Boot.
cmFORTH	Silicon Composers and others	Public-domain system for Novix Forth processor by C. Moore, ported to the Harris and SC-32 Forth processors by others.
Cyrano	Opto-22	Forth for a proprietary embedded controller
F-PC	T. Zimmer et al	Extensive public-domain system for the IBM-PC family
F83	Laxen and Perry	Public-domain system for the IBM-PC family, later ported by others to other platforms
HS/Forth	Harvard Softworks	IBM-PC family
JForth	Delta Research	Amiga
MacForth	Creative Solutions, Inc.	Apple Macintosh, NuBus interface boards
Mach2	Palo Alto Shipping	Apple Macintosh
mmsFORTH	Miller Microcomputer Services	IBM-PC family; business and commercial applications
MPEForth	MicroProcessor Engineering (UK)	PCs and embedded systems
mvpFORTH	Mountain View Press	Public-domain system on a variety of platforms
Open Boot	Sun Microsystems	Programmable ROM-based Forth on SPARC workstations
polyFORTH	FORTH, Inc.	Industrial systems on PCs and other platforms; interactive cross-compilers; consulting and custom programming services
URI/Forth	Laboratory Microsystems, Inc. (LMI)	IBM-PC family running DOS, OS2 and Windows; also cross compilers for a variety of systems

Table 3

Some major suppliers of Forth systems, services and related products.

In 1978, Maj. Selzer gave Don Colburn a copy of the 6502 Forth he wrote for Ragsdale in exchange for Colburn's writing two articles on Selzer's 6502 work. Colburn subsequently used this as a basis for a version based on the preliminary FORTH-77 standard (the only FORTH-77 implementation of which the authors are aware). In the Fall of 1979, Colburn generated a FIG-compatible system for prototypes of the 68000. A multitasking, multituser version of this product called MultiForth was demonstrated to Motorola in January, 1980, well ahead of production shipments of the 68000. When Hewlett Packard's desktop computer division designed a new generation of desktop computers around the 68000 in 1982, the first available third-party language product they distributed under an HP part number was MultiForth.

Colburn's company Creative Solutions also introduced MacForth, the first resident development system for the 128K Apple Macintosh, immediately after the Mac's debut in January of 1984. Because MacForth uniquely provided direct access to the entire Macintosh "Toolbox ROM" routines in a resident programming environment, along with comprehensive application examples, a majority of the first generation of Macintosh application programmers learned how to create and use pull-down menus, windows, graphics and mice with MacForth. Significant large-volume spreadsheets, 2D and 3D rendering and design packages, CAD/CAM design tools, games, medical diagnostics, image enhancement programs, accounting packages, desktop planetariums and process control applications were written on early Macintoshes in MacForth.

Byte Magazine dedicated its August, 1980 issue to Forth. It was their largest-selling issue to date, and was re-printed several times.

By 1985 there were over seventy vendors of Forth systems, ranging from single individuals to multi-million dollar organizations.

In 1982 Lawrence Forsley founded the Institute for Applied Forth Research, now called simply the Forth Institute. This organization sponsors an annual Conference on Forth Applications at the University of Rochester, Rochester NY, and publishes the *Journal of Forth Application and Research*, a refereed technical periodical on applications of Forth, new developments and techniques, and surveys of specific areas of Forth.

In 1989, George Shaw and others formed an ACM Special Interest Group on Forth called SIGForth, which also sponsors a newsletter and an annual conference.

3.2.1 Design principles

FIG Forth was optimized for portability rather than performance. Only a very few primitives were coded in assembler, and the rest of the logic was implemented using high-level Forth. As a result it was fairly slow — some operations, such as dictionary searches, were a factor of ten slower than representative commercial implementations.

Other internal decisions were similarly made with the neophyte in mind. For example, the earlier FORTH, Inc. systems compiled word names as the length of the name and the first three characters. This gave a lower collision rate than simple truncation, and was adequate most of the time. But the FIG model used variable-length names up to 31 characters, thereby trading size for user-friendliness. This was somewhat controversial at the time (see Fig. 1), but by the mid 1980's, most systems had converted to this usage.

The advent of personal computers provided Forth implementors with the incentive to learn to run under a host OS. The first non-native systems were

developed in 1980 by Martin Tracy of Micromotion (for the Apple II) and Ray Duncan of Laboratory Microsystems, Inc. (for CP/M on Z80's). LMI's system also featured a full-screen editor. In 1981, LMI added support for software and hardware floating point, and also pioneered performance enhancements such as native code translation and caching dictionary lookups in a hash table to accelerate dictionary searches.

```

DEA- EDI---
I AM AFR--- THA- THE LET--- IN THE LAS- ISS--- ABO---
FOR--- INC- VSI--- ONL- TER--- LET--- NAM- FIE--- BAS
HAD THE OPP----- EFF--- PRO- WEA- THE WAI--- MAN---
HIS LET--- ( LIE- THI- ONE ) SHO--- THA- SAV--- ONL-
TER--- LET--- AND COU--- IS JUS- ABO--- OPT--- IN
KEE--- OF A TRA--- OFF BET--- SAV--- MEM--- AND
WE STI--- DON--- SEE THE WEE- FOR 31 CHA----- NAM---
IN THE GEN----- CAS-
YOU--- TRU---
CHU--- MOO---
FOR--- INC-
    
```

Fig. 1
 "Letter to the Editor of Forth Dimensions [Moore, 1983] concerning the practice of storing names of Forth words as a count and first three characters.

The advent of non-native Forth implementations introduced an issue that remains controversial in Forth practice today, the use of host OS files for mass storage. There are two main approaches: abandoning traditional blocks altogether in favor of directly manipulating source and data in files, and mapping blocks to host OS files. The former approach is favored by implementors who are concentrating on systems for a particular OS (e.g., MS-DOS), whereas the latter is preferred by organizations such as FORTH, Inc. that support both native and non-native products.

Creative Solutions' MacForth used very compact object image strategies, including token threading and separated name heads to maximize the amount of memory available for program development on the original 128K Macintoshes. Other novel features included runtime relocation of the executable image and exclusion of word names in runtime systems without metacompilation. MacForth included a seamless programming environment incorporating screen based text editor, compiler, interpreter, and assembler in under 20k bytes of memory.

3.2.3 Influences

The FIG Model was in the public domain, and was ported to a wide variety of computer systems. Because the internal design of FIG Forth was essentially

the same across all machines, programs written in FIG Forth enjoyed a substantial degree of portability, even for "system-level" programs that directly manipulated the internals of dictionary entries and other implementation-dependent features. Because FIG Forth was the first introduction to Forth for many people, it is widely associated with "the nature of Forth."

However, FIG Forth was not representative of all commercial implementations of this era. Commercial vendors tended to be much more performance-conscious, and elected implementation strategies that optimized performance or size rather than porting ease, as we have seen.

The first major effort to standardize Forth was a meeting in Utrecht in 1977, attended by several astronomical Forth users and FORTH, Inc. (at that time the only commercial vendor). They produced a preliminary standard called FORTH-77, and agreed to meet the following year. Meetings in 1978 and 1979 on Catalina Island in California, now including representatives from the Forth Interest Group and other producers, yielded a more comprehensive standard called FORTH-79. Although FORTH-79 was very influential, many Forth users and vendors found flaws in it; in 1982 two meetings were held to update the standard, and in 1983 a new standard was released called FORTH-83. Both FORTH-79 and FORTH-83 specified a 16-bit, twos-complement, unaligned, linear byte-addressed virtual machine, and included a number of assumptions about implementation techniques.

Unfortunately, some of the changes in FORTH-83 produced grave incompatibilities with existing code. For example, the formal representation of a "true" flag had always been 1, and the word NOT inverted a Boolean flag. In FORTH-83, "true" became -1 and NOT became a bit-wise complement. Other problems involved the specification for floored division in FORTH-83 and a serious ambiguity in the specification of parameters for certain loop structures. The effect of these incompatibilities was divisive. Although most implementors agreed that FORTH-83 was an improvement and adopted the new standard, there remains a vocal group who never converted, and who remain skeptical of the whole standards process. Of the systems listed in Table 3, for example, most are fairly close to FORTH-83 compatibility; notable exceptions are MacForth, mmsFORTH and mypFORTH, all of which stayed with FORTH-79.

In 1981 Prentice Hall published *Starting FORTH*, by Leo Brodie [Brodie, 1981], then an employee of FORTH, Inc. Both lucid and entertaining (Brodie drew memorable cartoon figures representing important Forth primitives), *Starting FORTH* was also a thorough introduction to the language. It sold over 110,000 copies (for a time it was the best-seller in Prentice Hall's computer line) and exerted a powerful influence on many people learning about Forth for the first time, as well as on vendors scrambling to be compatible with it. Although the first edition was primarily based on FORTH, Inc.'s polyFORTH, it included many footnotes and examples in FIG Forth and other dialects. The second edition (1987) was based on the FORTH-83 standard.

Another major influence in the personal computer marketplace has been the competition between public-domain and commercial versions of Forth. In the mid-1980's, the FIG model was gradually replaced by the public domain F83 (produced by Henry Laxen, Mike Perry and others operating under the name "No Visible Support Software"), a multitasking system originally released on the IBM-PC. Versions have been developed by many independent programmers on a wide variety of other platforms. This system is so widespread that many people are led by its name to confuse it with the FORTH-83 standard. In fact, although it is largely compatible with FORTH-83, F83 goes well beyond the limited FORTH-83 standard in its features. In the late 80's, Tom Zimmer and others produced an even more extensive public-domain system for PCs called F-PC, which includes several megabytes of source code and utilities. But except for these, most public-domain Forths are rather limited.

Public-domain Forths have certainly helped to ensure that Forth is widely known. But their influence isn't entirely benign. According to Tyler Sperry, editor of *Embedded Systems Programming Magazine* [Sperry, 1991]:

The problem is that it is relatively easy to implement your own minimal Forth system. The kernel, after all, is only a few hundred bytes of code... Unfortunately, bringing up a Forth interpreter is like writing a Small C compiler: it's only a toy without a well-developed library. One of the biggest problems with public-domain and shareware systems is that their libraries are often only partially completed, with sketchy documentation. And that's putting the situation kindly.

People who have only seen or used limited public-domain Forth implementations often perceive that Forth itself is a toy. And suppliers of high-quality commercial systems must deal with prospective customers' assumptions that all Forths are the same, an assumption that naturally creates considerable price resistance given that the public-domain versions are extremely inexpensive. The standing joke within the Forth community, however, is that "when you've seen one Forth ... you've seen one Forth." The range in quality of code and documentation, nature and extent of libraries, as well as product support, is enormous. A prospective user is well advised to evaluate a number of both public domain and commercial offerings.

3.3 Embedded Systems

3.3.1 Environment and applications

Forth's ability to make maximum use of limited hardware resources made it a natural choice for embedded uses of microprocessors. Some of these have been small: an RCA 1802-based cardiac monitor (1979) that performed a detailed waveform analysis of heart beats was not much larger than the 1" x 2" tape cassette it used to record abnormalities. Some were large, such as the

750-ton stretch press used by Lockheed to form panels for the C5B airplane wings in the early 80's. Some were distributed, such as the roughly 500 networked processors used for an extensive facility management system at the King Khaled International Airport at Riyadh, Saudi Arabia [Rather, 1985]. Forth has been especially successful in developing firmware for hand-held devices made by companies such as Itron and MSI Data. In 1990, Federal Express won the prestigious Malcolm Baldrige quality award for its package-tracking system, whose data entry is performed by Forth-based hand-held devices carried by Federal's 50,000 couriers and agents world-wide.

Forth's extreme modularity facilitates thorough, systematic testing, which has made it attractive for applications requiring high reliability. As a result, it has been used in a number of satellites and Space Shuttle experiments. McDonnell Douglas used polyFORTH in their Electrophoresis in Space project [Wood, 1986] to control the cargo bay factory itself (multiple 68000 VME-bus boards), the astronaut's control console (a laptop PC), and their ground-based analysis computer (a Compaq PC). The November, 1990 Columbia shuttle flight carried four astronomy payloads, of which three were programmed in Forth [Ballard, 1991], and the January, 1992, Spacelab flight featured a Microgravity Vestibular Investigation (MVI) experiment using a polyFORTH system for on-board control and analysis [Paloski, 1986] and MACH2 in a ground-based Macintosh for analysis.

Probably the most prolific single purveyor of embedded Forths is Sun Microsystems, whose SPARC workstations all use a programmable Forth-based monitor called Open Boot, developed by Mitch Bradley and associates. Bradley believes [Bradley, 1991] that Forth was successful for this purpose because it offered:

1. a CPU-independent "virtual machine" to use for the byte-coded portable drivers;
2. a debugging environment for those drivers;
3. an interactive command language, with complete programming language capability, that was useful for hardware startup and debugging;
4. a built-in debugging environment for the firmware itself (firmware is otherwise rather painful to debug);
5. a debugging environment for the operating system software;
6. extensibility, allowing easy support of new hardware requirements and features; and
7. great flexibility in tuning the implementation for speed/space tradeoffs.

At least one other major board-level CPU vendor has adopted Open Boot firmware across their product line, and there is a working group developing an IEEE standard for it.

3.3.2. Design principles

From the minis of the '70s to the PCs of the '80s, most Forth systems have supported development on the same computer on which the completed application is to run. Even the microprocessor systems of the late '70s and early '80s were developed on the same CPU (as opposed to cross-development), with development software features for stripping the development tools and producing a ROMable target.

Most embedded systems lack a disk, a terminal, or both, thereby rendering themselves inhospitable to even the leanest Forth programming environment. Nonetheless, some vendors do provide on-board Forths in microcontrollers. Examples include the Rockwell AIM 65 mentioned above, and microcontroller boards sold by New Micros, Inc. of Texas; Vesta Technologies, Inc., in Colorado; and Opto-22 in California.

But as PCs became ubiquitous, they also became popular as hosts for more comfortable and powerful Forth cross-development environments. These have generally been based on modified versions of the classical Forth meta-compilers, adapted to support cross development.

The traditional Forth dictionary is integrated: a "definition" includes the word's name (which can be found in a dictionary search performed by the text interpreter), an executable portion (typically a pointer to code which executes words of a particular class, such as colon definitions, variables, constants, etc.), and data space (containing one or more values or addresses of words that make up the content of the definition), all classically in contiguous memory locations (but see Section 5.2, Implementation Strategies). A metacompiler divides these structurally into portions that are used by the host system's compiler (equivalent to a symbol table) and portions required at run-time in the target. In order for a target program to be ROMable, the compiler must also manage separate ROM and RAM data spaces, usually using multiple sets of dictionary pointers.

4. Hardware Implementations of Forth

The internal architecture of Forth simulates a computer with two stacks, a set of registers, and other well-defined features. As a result, it was almost inevitable that someone would attempt to build a hardware representation of the actual Forth computer.

The first such effort was made in 1973 by John Davies, of the Jodrell Bank Radio Astronomy Observatory near Manchester, England. Davies' approach was to re-design a Ferranti computer that had gone out of production to optimize its instruction set for Forth.

The first actual Forth computers were bit-sliced board-level products. The first of these was made by a California company called Standard Logic, in 1976.

By making a minor modification in the instruction set of their board-level computer, Standard Logic's chief programmer Dean Sanderson was able to implement the precise instruction that Forth uses in its "address interpreter" to move from one high-level command to the next. Their system was used widely by the U. S. Post Office.

In the early 1980's, Rockwell produced a microprocessor with Forth primitives in on-chip ROM, the Rockwell AIM 65F11 [Dumse, 1984]. This chip has been used quite successfully in embedded microprocessor applications. However, no attempt was made to adapt the actual architecture of the processor (basically a 6502) for Forth support.

In 1981, Moore himself undertook to design an actual Forth chip. Working first at FORTH, Inc. and subsequently with a start-up company called Novix, formed to develop the chip, Moore completed the design in 1984, and the first prototypes were produced in early 1985 [Golden, 1985]. This design was subsequently purchased and adapted by Harris Semiconductor Corp., and formed the basis of their line of RTX processors.

Starting in the early 1980's, a group at the Johns Hopkins Applied Physics Laboratory in Maryland developed a series of experimental Forth processors for use in space instrumentation [Hayes, 1987]. The most successful of these, marketed as the SC-32 by Silicon Composers of Palo Alto, CA, was used to control the Hopkins Ultraviolet Telescope which flew in the Columbia Space Shuttle in November, 1990 [Ballard, 1991]. It continues to be the basis for more space instruments under development.

Moore himself, working on his own, has continued to develop Forth-based processors for special applications.

The various Forth processors have had an influence on Forth software systems. In order to take full advantage of these architectures, Forth compilers were developed by Moore, FORTH, Inc. and Laboratory Microsystems that generated machine code optimized for the chip's internal architecture. A native looping structure in the Novix and Harris chips called FOR ... NEXT (which counted down from a single-argument upper limit to zero) led to adoption of this structure in other Forths as well.

5. Present and Future Directions

The computer industry has always been characterized by rapid and profound changes. Since Forth was last standardized in 1982-3, the speed, memory size and disk capacity of affordable personal computers have increased by factors of more than 100. 8-bit processors are now rare in PCs (although they are still widely used in embedded systems), and 32-bit processors are common. Operating systems, programming environments and user interfaces are far

more sophisticated. Many recent Forth implementations, both commercial and public-domain, have attempted to address these issues.

5.1 Standardization Efforts

At the time of writing (November, 1992), a Technical Committee X3J14 (of which authors Rather and Colburn are members) is nearing completion of an ANS Forth. Among the 20 voting members in the TC are vendors (FORTH, Inc., Creative Solutions, Sun Microsystems and a division of NCR), some large user organizations (Ford Motor Co., NASA), and a number of smaller user organizations, consultants and experts. Starting in 1987, this group has addressed a number of problems with FORTH-79 and FORTH-83, as well as some contemporary issues. A few of the issues addressed in the draft standard follow, as they represent current areas of lively debate and technical activity among Forth users and implementors.

ANS Forth attempts to reconcile some of the divisions caused by the incompatibilities between FORTH-79 and FORTH-83. For example, it retains `0=` to perform the FORTH-79 `NOT` function, introduces `INVERT` to perform the FORTH-83 `NOT`, and removes the word `NOT`. This enables application writers who depend on either version to leave their programs unchanged, and achieve compatibility by adding a simple shell in which `NOT` is defined as a synonym for the preferred behavior.

The proposed standard also removes virtually all restrictions on implementation options, provides for independence from CPU word size, and offers a number of optional extension word-sets for functions such as host OS file compatibility, dynamic memory allocation and floating point arithmetic. Some significant issues addressed by ANS Forth follow.

5.1.1 Cell size

FORTH-79 and FORTH-83 mandated a 16-bit architecture, including stack width, addresses, flags and numbers. ANS Forth specifies sizes in terms of a "cell," whose width is implementation-defined but must be at least 16 bits. Words have been added to increment addresses transportably by a cell, a character, or an integral number of cells or characters.

5.1.2 Arithmetic

Amid great controversy, FORTH-83 mandated floored division. Not only was this incompatible with prior usage (which didn't specify the algorithm for handling signed division), it is also at variance with hardware multiply/divide instructions on most processors. But many people felt strongly that floored division is mathematically more appropriate, and that it was important to specify. Recognizing that there were many implementations on both sides of this issue, the TC opted to allow either floored or truncated

division. The implementation must specify which default it uses, and must provide primitives supporting both methods.

5.1.3 Control structures

One of the unique characteristics of Forth is the degree to which its own internal tools are accessible to the application programmer. For example, there is one lexical analyzer used by the compiler, assembler, and text interpreter; it is also available for command and text parsing in applications. Similarly, the tools that implement control structures such as loops and conditionals are available for making custom structure words. In 1986 Will Baden demonstrated [Baden, 1986] that the standard Forth structure words plus a few extensions made from these underlying tools are adequate to make any structure, including solutions to problems posed in D. E. Knuth's paper "Structured Programming with `go` to statements" [Knuth, 1974].

Structure	Description
<code>DO ... LOOP</code>	Finite loop incrementing by 1
<code>DO ... <n> +LOOP</code>	Finite loop incrementing by <n>
<code>BEGIN ... <<> UNTIL</code>	Indefinite loop terminating when <<> is 'true'
<code>BEGIN ... <<> WHILE ... REPEAT</code>	Indefinite loop terminating when <<> is 'false'
<code>BEGIN ... AGAIN</code>	Infinite loop
<code><<> IF ... ELSE ... THEN</code>	Two-branch conditional; performs words following <code>IF</code> if <<> is 'true' and words following <code>ELSE</code> if it is 'false'. <code>THEN</code> marks the point at which the paths merge.
<code><<f> IF ... THEN</code>	Like the two-branch conditional, but with only a 'true' clause.

Table 4

Standard control structures in Forth. ANS Forth allows programmers to form new structures by mixing the component words or using them to define new structure words.

FORTH-79 and FORTH-83 provided syntactic specifications for the common structure primitives. The latter were not widely adopted, however, and few implementations perform the kind of syntax checking the standards anticipated. F83 offers a limited form of syntax checking, in that it requires the stack (which is used at compile time for compiling structures) to have the same size before and after compiling a definition, the theory being that a stack imbalance would indicate an incomplete structure. Unfortunately, this technique prevents the very common practice of leaving a value on the compile-time stack which is to be compiled as a literal inside a definition.

Common practice often took advantage of knowledge about how the structure words worked at compile time to manipulate them in creative ways.

The ANS Forth Technical Committee sanctioned this by providing specifications of both the compile-time and run-time behaviors of the structure words, so that they may be combined in arbitrary order. A set of structure primitives is provided in a "programming tools" wordset, and the word **POSTPONE** is provided to enable programmers to write new structure words that reference existing compiler directives in order to provide a portion of the desired new behavior.

5.2 Implementation Strategies

The original Forth systems developed by Moore in the 1970's compiled source from disk into an executable form in memory. This avoided the separate compile-link-load sequences characteristic of most compiled languages, and led to a very interactive programming style in which the programmer could use the resident Forth editor to modify source and re-compile it, having it available for testing in seconds. The internal structure of a definition was as shown in Fig. 2, with all fields contiguous in memory. The FIG model and its derivatives modified the details of this structure somewhat, but preserved its essential character.

Forth systems implemented according to this model built a high-level definition by compiling pointers to previously defined words into its parameter field; the address interpreter that executed such definitions proceeded through these routines, executing the referenced definitions in turn by performing indirect jumps through the register used to keep its place. This is generally referred to as *indirect-threaded* code.




Fig. 2

Diagram showing the logical components of a Forth definition. In classical implementations, these fields are contiguous in memory. The data field will hold values for data objects, addresses or tokens for procedures, and the actual code for code definitions.

The need to optimize for different conditions has led to a number of variants in this basic implementation strategy however. Some of the most interesting are:

1. **Direct threaded code.** In this model, the code field contains machine code instead of a pointer to machine code. This is somewhat faster, but typically costs extra bytes for some classes of words. It is most prevalent on 32-bit systems.

2. **Subroutine-threaded code.** In this model, the compiler places a jump-to-subroutine instruction with the destination address in-line. This technique costs extra bytes for each compiled reference on a 16-bit system. It is often slower than direct-threaded code, but it is an enabling technique to allow the progression to native code generation.
2. **Native code generation.** Going one step beyond subroutine-threaded code, this technique generates in-line machine instructions for simple primitives such as + and jumps to other high-level routines. The result can run much faster, at some cost in size and compiler complexity. Native code can be more difficult to debug than threaded code. This technique is characteristic of optimized systems for the Forth chips such as the RTX, and on 32-bit systems where code compactness is often less critical than speed.
3. **Optimizing compilers.** A variant of native code generation, these were invented for the Forth processors (discussed in Section 1.5) that can execute several Forth primitives in a single cycle. They looked for the patterns that could be handled in this way and automatically generated the appropriate instruction. The range of optimization was governed by the capabilities of the processor; for example, the polyFORTH compiler for the Novix and RTX processors had a four-element peephole window.
4. **Token threading.** This technique compiles references to other words using a token, such as an index into a table, which is more compact than an absolute address. Token threading was used in a version of Forth for a Panasonic hand-held computer developed in the early 1980's, for example, and is a key element in MacForth.
5. **Segmented architectures.** The 80x86 family supports segmented address spaces. Some Forths take advantage of this to enable a 16-bit system to support programs larger than 64K. Similarly, implementations for Harvard-architecture processors such as the 8051 and TI TMS320 series manage separate code and data spaces.

Although the early standards assumed the classical structure, ANS Forth makes a special effort to avoid assumptions about implementation techniques, resulting in prohibitions against assuming a relationship between the head and data space of a definition or accessing the body of a data structure other than by pre-defined operators. This has generated some controversy among programmers who prefer the freedom to make such assumptions over the optimizations possible with alternative implementation strategies.

5.3 Object-oriented Extensions

Forth's support for custom data types with user-defined structure as well as compile-time and run-time behaviors has, over the years, led programmers

to develop object-based systems such as Moore's approach to image processing described above (Section 2.2.2, item 2). Pountain [1987] described one approach to object-oriented programming in Forth, which has been tried by a number of implementors. Several Forth vendors have taken other approaches to implementing object-based systems, and this is currently one of the most fertile areas of exploration in Forth.

In 1984, Charles Duff introduced an object-oriented system written in Forth called Neon [Duff, 1984 a, b]. When Duff discontinued supporting it in the late 80's it was taken over by Bob Lowenstein, of the University of Chicago's Yerkes Observatory, where it is available as a public-domain system under the name Yerk. More recently, Michael Hore re-implemented Neon using a subroutine-threaded code; the result is available (also in the public domain) under the name MOPS. Both Yerk and MOPS are available as down-loadable files on a number of Forth-oriented electronic bulletin boards listed at the end of this paper.

6. A Posteriori Evaluation

The early development of FORTH was in many ways quite different from that of most other programming languages. Whereas they generally emerged full featured with unambiguous formal specifications for language syntax and semantics, Forth enjoyed a lengthy, dynamic adolescence, in which each fundamental presupposition of the language was tested on the anvil of actual applications experience. During this period, Moore, unencumbered by a large following of users, often made revolutionary changes to the language on a daily basis to suit his current view of what the language should be. He had complete control and responsibility for the machine at hand, from the first bootstrap loader to the completed application. The language converged toward the actual needs of one man solving a broad class of technically challenging problems in resource-constrained environments.

The resulting method of problem solving, expressed by the resulting *de facto* language specification, has proven useful to others. Given the complete flexibility to add syntax checking, data typing, and other more formal structures often considered essential to programming languages, most of the several hundred people who have independently implemented versions of Forth for their own use have not done so. The results of their efforts, as surveyed by the ANS Forth Technical Committee, represent a startlingly democratic ratification of Moore's personal vision.

6.1 Meeting Objectives

Without a formal language design specification citing clearly defined objectives, we can only evaluate the stated objectives of the inventor of the language and those who have used it. Personal productivity and intellectual portability were Moore's primary stated objectives. Forth has been ported

across the vast majority of programmable computers and has been embodied in several different dedicated Forth computer architectures.

In 1979, Chuck Moore looked back on ten years' experience with Forth and observed [Moore, 1979]:

My original goal was to write more than 40 programs in my life. I think I have increased my throughput by a factor of 10. I don't think that throughput is program-language limited any longer, so I have accomplished what I set out to do. I have a tool that is very effective in my hands — it seems that it is very effective in others' hands as well. I am happy and proud that this is true.

Today he sees no reason to change this assessment.

The developers of FIG Forth saw their systems spread all over the world, along with chapters of their organization, and influence Forth programmers everywhere. Their goal of instantiating additional commercial vendors of Forth products was also achieved.

Of the many entrepreneurs who committed their careers and fortunes to Forth-based enterprises, few have become rich and famous for their efforts. But most have had the satisfaction of seeing their own productivity increased just as Moore did, and of having seen seemingly impossible project objectives met due to the power and flexibility of the language. They have also enjoyed prosperity in making this capability available to their clients and customers.

Given Moore's criteria of productivity and portability, perhaps the best measure of achieving these objectives is the very large quantity and range of application programs that have been written in Forth by a small number of programmers across a very broad variety of computers.

6.2 Major Contributions of Forth

In 1984, Leo Brodie wrote a book on designing Forth applications called *Thinking Forth* [Brodie, 1984]. In it, he quoted a number of Forth programmers on their design and coding practices. In an Epilogue, several of them commented that Forth had significantly influenced their programming style in other languages, and indeed their approaches to problem solving in general. Here are two examples, which are typical of observations of Forth users in general:

[The] essence of good Forth programming is the art of factoring procedures into useful free-standing words. The idea of the Forth word had unexpected implications for laboratory hardware design.

Instead of building a big, monolithic, all-purpose Interface, I found myself building piles of simple little boxes which worked a lot like Forth words: they had a fixed set of standard inputs and outputs, they performed just one function, they were designed to connect up to each

other without much effort, and they were simple enough that you could tell what a box did just by looking at its label....

Because Forth is small, and because Forth gives its users control over their machines, Forth lets humans control their applications. It's just silly to expect scientists to sit in front of a lab computer playing "twenty questions" with packaged software. Forth ... lets a scientist instruct the computer instead of letting the computer instruct the scientist.

— Mark Bernstein, president of Eastgate Systems, Inc., Cambridge, MA

Forth has changed my thinking in many ways. Since learning Forth I've coded in other languages, including assembler, BASIC and FORTRAN. I've found that I used the same kind of decomposition we do in Forth, in the sense of creating words and grouping them together.

More fundamentally, Forth has reaffirmed my faith in simplicity. Most people go out and attack problems with complicated tools. But simpler tools are available and more useful.

— Jerry Boutelle, owner of Nautilus Systems, Santa Cruz, CA

Mitch Bradley reports [Bradley, 1991] that the design of the Forth-based Open Boot has significantly influenced the thinking of the people at Sun Microsystems who are responsible for the low-level interfaces in the Unix kernel. Open Boot design philosophy is influencing driver interfaces, the device naming system, and the early startup and configuration mechanisms. There is even talk of unifying the syntax of several disparate kernel configuration files by using Forth syntax and including a subset Forth interpreter in the Unix kernel. People at Sun who have worked with Open Boot are impressed by the fact that the simple postfix syntax never "runs out of steam" or "paints you into a corner."

6.3 Mistakes or Desired Changes

Forth has a chameleon-like capacity to adapt to particular application needs. Indeed, the process of programming in Forth is to add to it application-oriented words at increasingly high levels until all the desired functionality is implemented. So for any project or even any particular programming group, any perceived needs will be promptly addressed. When looking for "mistakes," then, the most useful questions to ask are what were the things that a significant number of implementors have chosen to change or add, and what characteristics of the language may have prevented its wider acceptance.

One of the first actions taken by the ANS Forth Technical Committee when it formed in 1987 was to poll several hundred Forth implementors and users to determine their views on problems in the language that needed to be addressed. The issues cited fell into three categories: "mistakes" in one or both of the existing standards (e.g., incompatibilities introduced by FORTH-83 and anomalies such as an awkward specification for arguments to DO); obsolete

restrictions in FORTH-83 (mainly the reliance on a 16-bit architecture), and a need for standards for such things as host file access, floating point arithmetic, etc. Features in the latter group were by then offered by most commercial and many public-domain systems, but as they had been developed independently there was variance in usage and practice. ANS Forth has attempted to address all these concerns.

In retrospect, however, the lack of standard facilities for such things as floating point arithmetic, which are covered by other languages, has probably impeded widespread acceptance of Forth. It's insufficient to point out that most commercial systems offer them if the public perception of the language is formed by a standard that omits any mention of such features! From this perspective, the ANS Forth effort has come almost too late.

Another difficulty is that Forth's very identity is unclear: it is not only unconventional in appearance, with its reliance on an overt stack architecture and postfix notation, but it broadly straddles territory conventionally occupied by not only languages but also operating systems, editors, utilities, etc., that most people are accustomed to viewing as independent entities. As a result, it's difficult to give a simple answer to the question of what it is.

The integrated character of Forth is viewed by its practitioners as its greatest asset. As Bradley [1991] expresses it,

Forth has taught me that 'firewalls' between different components of a programming environment (i.e., the different syntax used by compilers, linkers, command interpreters, etc.) are very annoying, and it is much more pleasant to have a uniform environment where you can do any thing at any level at any time, using the same syntax.

Duncan [1991], however, believes that this seamless integration of Forth the language, Forth the virtual machine and Forth the programming environment is a significant barrier to mainstream acceptance. He notes that the same has been observed regarding Smalltalk vs. C++:

I have been using C++ for some months now, and the very things about C++ that frustrate me — the language is not written in itself (thus there is no way to use the building blocks for the programming environment as part of the application), the language is not truly extensible (e.g. the operators for the native data types cannot be overridden), and there is no programming environment that is smart about the language and class hierarchies — are the things that traditional language experts see as assets for C++ compared to Smalltalk!

So long as Forth users are convinced that its integrated, intrinsically inter-active character is the key to their productivity as programmers, however, it is unlikely to change.

6.4 Problems

Some languages tend to be "levelers:" that is, a program written by an expert is unlikely to be significantly better (smaller, faster, etc.) than one written by a novice. Chuck Moore once observed [Moore, 1979], "...FORTH is an amplifier. A good programmer can do a fantastic job with FORTH; a bad programmer can do a disastrous one." While never quantified, this observation has been repeated on many Forth projects across a broad programmer population, and has achieved the status of "folk wisdom" within the Forth community.

This tendency has given Forth the reputation of being "unmanageable," and there have been some highly publicized "Forth disasters" (notably Epson's VALDOCS project in the early 1980's). On close examination, however, the root causes of this and other failed Forth projects are the same problems that doom projects using other languages: inadequate definition, poor management and unrealistic expectations.

There have also been a number of Forth successes, such as the facility management system for the Saudi Arabian airport mentioned above, in which a project that was estimated to contain 300,000 lines of executable Fortran, PLM and assembly language software was totally redesigned, recoded in Forth and tested to the satisfaction of the customer in only 18 months [Rather, 1985]. The result ran more than a factor of 10 faster.

Jack Woehr, a senior project manager for Vestia Technologies, observes [Woehr, 1991] that successful management of Forth projects demands nothing more than generally good management practices plus an appreciation of the special pride that Forth programmers take in their unusual productivity. Forth rewards a management style that believes a small team of highly skilled professionals can do a better job, in a shorter time, at less overall cost than a large group of more junior programmers.

6.5 Implications for Current and Future Languages

What can be learned from twenty years' experience with Forth? Forth stands as a living challenge to many of the assumptions guiding language developers. Its lack of rigid syntax and strong data typing, for example, are characteristically listed as major advantages by Forth programmers. The informal, interactive relationship between a Forth system and its programmer has been shown through many projects to shorten development times in comparison with more conventional tools such as C. Despite the tremendous increases in the size and power of modern computers, the Forth's combination of easy programming, compact size and fast performance (characteristics often thought to be mutually exclusive) continues to earn a loyal following among software developers, especially for embedded systems.

References

- [ANS 1991] *Draft Proposed ANS Forth*, document number X3.215-199x, available from Global Engineering Documents, 2805 McGaw Ave., Irvine, CA, 92714.
- [Baden, 1986] Baden, W., "Hacking Forth." *Proceedings of the Eighth FORML Conference*, pub. by the Forth Interest Group, P. O. Box 8231, San Jose, CA 95155, 1986.
- [Ballard, 1991] Ballard, B. and Hayes, J., "Forth and Space at the Applied Physics Laboratory," in *Proceedings of the 1991 Rochester Forth Conference*. Rochester, NY: Forth Institute, 1991.
- [Bradley, 1991] Bradley, M., private communication, 7/8/91.
- [Brodie, 1981] Brodie, L. *Starting FORTH*. Englewood Cliffs, NJ: Prentice Hall, 1981.
- [Brodie, 1984] Brodie, L. *Thinking FORTH*. Englewood Cliffs, NJ: Prentice Hall, 1984.
- [Cox, 1987] William C. Cox, "A case for NPOs in real-time applications." *I&CS Magazine* (pub. by Chilton), October, 1987.
- [Derrick, 1982] Derrick, M. and Baker, L., *The Forth Encyclopedia*. Mountain View, CA: The Mountain View Press, 1982.
- [Dewar, 1970] Dewar, R., "Indirect Threaded Code." *Communications of the ACM*, **18**, 6, 1975.
- [Dijkstra, E. W. 1970, "Structured Programming," *Software Engineering Techniques*, Buxton, J. N., and Randell, B., eds. Brussels, Belgium, NATO Science Committee, 1969.
- [Duff, 1984a] C. Duff and N. Iverson, "Forth Meets Smalltalk," *Journal of Forth Application and Research*, **2**, 1, 1984.
- [Duff, 1984b] C. Duff, "Neon — Extending Forth in New Directions," *Proceedings of the 1984 Asilomar FORML Conference* pub. by the Forth Interest Group, P. O. Box 8231, San Jose, CA 95155, 1984.

- [Dumse, 1984]
Dumse, R., "The R65F11 and F68K Single Chip FORTH Computers," *Journal of Forth Application and Research*, 2, 1, 1984.
- [Duncan, 1988]
Duncan, R. (Gen'l Ed.). *The MS-DOS Encyclopedia*. Redmond, WA: Microsoft Press, 1988.
- [Duncan, 1991]
Duncan, R., private communication, 7/5/91.
- [Electronics, 1976]
"RCA may offer memory-saving processor language." *Electronics* (Feb. 19, 1976, p. 26).
- [FIC, 1978]
Forth Dimensions, 1, 1, June/July 1978, pub. by the Forth Interest Group, P. O. Box 8231, San Jose, CA 95155.
- [FIC, 1983]
Forth Dimensions, 5, 3, September/October, 1983, pub. by the Forth Interest Group, P. O. Box 8231, San Jose, CA 95155.
- [Gehrz, 1978]
Gehrz, R. D., and Hackwell, J. A., "Exploring the Infrared Universe from Wyoming." *Sky and Telescope* (June, 1978).
- [Golden, 1985]
Golden, J., Moore, C. H. and Brodie, L. "Fast Processor Chip Takes Its Instructions Directly from Forth," *Electronic Design* (March 21, 1985).
- [Hart, 1968]
Hart, J. F. et al., *Computer Approximations*. Malabar, FL: Krieger, 1968; 2nd ed., 1978.
- [Hayes, 1987]
Hayes, J. R., Fraeman, M. E., Williams, R. L. and Zarella, T., "A 32-bit Forth Microprocessor," *Journal of Forth Application and Research*, 5, 1, 1987.
- [Knuth, 1974]
Knuth, D. E., "Structured Programming with go to statements." *Computing Reviews*, 1974, #4.
- [Moore, 1958]
Moore, Charles H. and Lautman, D. A., "Predictions for Photographic Tracking Stations — APO Ephemeris 4" in *SAO Special Report #11*, G. F. Schilling, ed. Cambridge, MA: Smithsonian Astrophysical Observatory, 1958.
- [Moore, 1970a]
Moore, Charles H. and Leach, G. C. *FORTH — A Language for Interactive Computing*. Amsterdam, NY: Mohasco Industries Inc. (internal pub.) 1970.

- [Moore, 1970b]
Moore, C. H., *Programming a Problem-oriented Language*. Amsterdam, NY: Mohasco Industries Inc. (internal pub.) 1970.
- [Moore, 1974a]
Moore, C. H., "FORTH: A New Way to Program a Computer," *Astronomy & Astrophysics Supplement Series*, 15, 3, June 1974. Proceedings of the Symposium on Collection and Analysis of Astrophysical Data at NRAO, Charlottesville, VA, Nov. 13-15, 1972.
- [Moore, 1974b]
Moore, C. H. and Rather, E. D., "The FORTH Program for Spectral Line Observing on NRAO's 36 ft Telescope" *Astronomy & Astrophysics Supplement Series*, 15, 3, June 1974. This is the Proceedings of the Symposium on the Collection and Analysis of Astrophysical Data given at NRAO, Charlottesville, VA, November 13-15, 1972.
- [Moore, 1979]
Moore, C. H., "FORTH, The Last Ten Years and the Next Two Weeks..." Address at the first FORTH Convention, San Francisco, CA, October 1979, reprinted in *Forth Dimensions*, 1, 6, 1980.
- [Moore, 1983]
Letter to the Editor of *Forth Dimensions*, 3, 1, 1983.
- [Paloski, 1986]
Paloski, W. H., Odette, L., and Krever, A. J., "Use of a Forth-based Prolog for Real-time Expert System." *Journal of Forth Application and Research*, 4, 2, 1986.
- [Pountain, 1987]
Pountain, R. *Object Oriented Forth*. New York: Academic Press, 1987.
- [Parnas, 1971]
Parnas, D. L., "Information Distribution Aspects of Design Methodology." *Proc. IFIP 1971 Congress*. Ljubljana, Yugoslavia. (Phys. Sci. 1975)
- "Graphics in Kitt Form." *Physical Science*, Nov. 1975, p. 10.
- [Rather, 1972]
Rather, E. D. and Moore, C. H., *FORTH Programmer's Guide*, NRAO Computer Division Internal Report #11, 1972. A later version, with J. M. Hollis added as a co-author, was Internal Report #17, 1974.
- [Rather, 1976a]
Rather, E. D. and Moore, C. H., "The FORTH Approach to Operating Systems," *Proceedings of the ACM*, Oct. 1976 pp. 233-240.
- [Rather, 1976b]
Rather, E. D. and Moore, C. H., "High-level Programming for Microprocessors", *Proceedings of Electiro 76*.

[Rather, 1985]
 Rather, E. D., "Fifteen Programmers, 400 Computers, 36,000 Sensors and Forth," *Journal of Forth Application and Research* (3, #2, 1985). Available from the Forth Institute, P. O. Box 27686, Rochester, NY, 14627.

[Sperry, 1991]
 Sperry, Tyler, "An Enemy of the People." *Embedded Systems Programming* (4, 12), December, 1991.

[Taylor, 1980]
 Taylor, Alan, "Alternative Software Making Great Strides." *Computerworld*, 12/?/80.

[TRS-80, 1979]
 Press release published in "Software and Peripherals" section of *Micomputer News*, 8/30/79.

[Veis, 1960]
 Veis, George and Moore, C. H., "SAO Differential Orbit Improvement Program" in *Tracking Programs and Orbit Determination Seminar Proceedings*. Pasadena CA: Jet Propulsion Laboratories, 1960.

[Woehr, 1991]
 Woehr, Jack J., "Managing Forth Projects," *Embedded Systems Programming* (May, 1991).

[Wood, 1986]
 Wood, R. J., "Developing Real-time Process Control in Space." *Journal of Forth Application and Research*, 4, 2, 1986.

Bibliography
 Brodie, L. *Starting FORTH*. Englewood Cliffs, NJ: Prentice Hall, 1981.
 Brodie, L. *Thinking FORTH*. Englewood Cliffs, NJ: Prentice Hall, 1984.
 Feierbach, G. and Thomas, P. *Forth Tools & Applications*. Reston, VA: Reston Computer Books, 1985.
 Haydon, G. B. *All about Forth: An Annotated Glossary*. La Honda CA: Mountain View Press, 1990.
 Kelly, M. G., and Spies, N. *FORTH: A Text and Reference*. Englewood Cliffs, NJ: Prentice Hall, 1986.
 Knecht, K. *Introduction to Forth*. Howard Sams & Co., Indiana, 1982.
 Kogge, P. M. "An Architectural Trail to Threaded Code Systems." *IEEE Computer* (March, 1982).
 Koodman, P. *Stack Computers. The New Wave*. Chichester, West Sussex, England. Ellis Horwood Ltd. 1989

Martin, T. *A Bibliography of Forth References*, 3rd Ed. Rochester, NY: Institute for Applied Forth Research, 1987.

McCabe, C. K. *Forth Fundamentals* (2 volumes). Oregon: Dilithium Press, 1983.

Moore, C. H. "The Evolution of FORTH — An Unusual Language." *Byte* (August 1980).

Ouverson, Marlin (ed). *Dr. Dobbs Toolbook of Forth*. Redwood City, CA: M&T Press, Vol. 1, 1986 Vol 2, 1987.

Pountain, R. *Object Oriented Forth*. New York: Academic Press, 1987.

Rather, E. D. "Forth Programming Language." *Encyclopedia of Physical Science & Technology* (Vol. 5). New York: Academic Press, 1987.

Rather, E. D. "FORTH." *Computer Programming Management*. Auerbach Publishers, Inc., 1985.

Terry, J. D. *Library of Forth Routines and Utilities*. New York: Shadow Lawn Press, 1986

Tracy, M. and Anderson, A. *Mastering Forth* (2nd ed). New York: Brady Books, 1989.

Winfield, A. *The Complete Forth*. New York: Wiley Books, 1983.

Forth Organizations:
 ACM/SIGForth, Association for Computing Machinery, 11 W. 42nd St., New York, NY 10036. Publishers of the quarterly SIGForth newsletter, sponsor of annual conference (usually in March).
 The Forth Institute, 70 Elmwood Dr., Rochester, NY 14611. Publishers of the *Journal of Forth Application and Research* and sponsors of the annual Rochester Forth Conference (usually in June). *Proceedings* of these Conferences are also published.
 The Forth Interest Group, P. O. Box 8231, San Jose, CA 95155. Publishers of *Forth Dimensions* (newsletter), source of various books on Forth and public-domain Forth systems; also coordinator for worldwide chapters and sponsor of annual FORML conference (usually in November).
On-line Resources:
 BIX (ByteNet) (for information call 800-227-2983) Forth Conference. Access BIX via TymNet, then type j Forth. Type FORTH at the : prompt.
 CompuServe (for information call 800-848-8990) Forth Forum sponsored by Creative Solutions, Inc. Type GO FORTH at the ! prompt.
 GEnie (for information call 800-639-9636) Forth Round Table. Call GEnie local node, then type FORTH. Sponsored by the Forth Interest Group. Also connected via cross-postings with Internet (comp.lang.forth) and other Forth conferences.