

The History of Fortran I, II, and III

JOHN BACKUS

This article discusses attitudes about “automatic programming,” the economics of programming, and existing programming systems, all in the early 1950s. It describes the formation of the Fortran group, its knowledge of existing systems, its plans for Fortran, and the development of the language in 1954. It describes the development of the optimizing compiler for Fortran I, of various language manuals, and of Fortran II and III. It concludes with remarks about later developments and the impact of Fortran and its successors on programming today.

[Editor’s note: This paper originally appeared in the *Annals of the History of Computing*, vol. 1, no. 1, July 1979. The sayings used as “pull quotes” were not in the original paper but were on cards that were distributed during the 1982 AFIPS National Computing Conference, which included a “Pioneer Day” honoring the Fortran pioneers. Readers interested in the topic of Fortran should consult the *Annals*, vol. 6, no. 1, which was a special issue devoted to the history of the language during its first 25 years.]

Early Background and Environment Attitudes About Automatic Programming in the 1950s

Before 1954, almost all programming was done in machine language or assembly language. Programmers rightly regarded their work as a complex, creative art that required human inventiveness to produce an efficient program. Much of their effort was devoted to overcoming the difficulties created by the computers of that era:

- the lack of index registers,
- the lack of built-in floating-point operations,
- restricted instruction sets (which might have AND but not OR, for example), and
- primitive input-output arrangements.

Given the nature of computers, the services that “automatic programming” performed for the programmer were concerned with overcoming the machine’s shortcomings. Thus, the primary concern of some “automatic programming” systems was to allow the use of symbolic addresses and decimal numbers (e.g., the MIDAC Input Translation program¹⁸).

But most of the larger “automatic programming” systems (with the exception of J.H. Laning and N. Zierler’s algebraic system²⁷ and the A-2 compiler^{29,36}) simply provided a synthetic “computer” with an order code different from that of the real machine. This synthetic computer usually had floating-point instructions and index registers and had improved input-output commands; it was, therefore, much easier to program than its real counterpart.

The A-2 compiler also came to be a synthetic computer sometime after early 1954. But in early 1954, its input had a much cruder form; instead of “pseudoinstructions,” its input was then a

complex sequence of “compiling instructions” that could take a variety of forms, ranging from machine code itself to lengthy groups of words constituting rather clumsy calling sequences for the desired floating-point subroutine, to “abbreviated form” instructions that were converted by a “Translator” into ordinary “compiling instructions.”²⁹

After May 1954, the A-2 compiler acquired a “pseudocode” that was similar to the order codes for many floating-point interpretive systems that were already in operation in 1953, for example:

- the Los Alamos systems DUAL and SHACO,^{16,40}
- the Massachusetts Institute of Technology (MIT) “Summer Session Computer,”¹
- a system for the Illiac designed by D.J. Wheeler,³⁰ and
- the Speedcoding system for the IBM 701.⁵

The Laning and Zierler system was quite a different story: It was the world’s first operating algebraic compiler, a rather elegant but simple one. Donald Knuth and Luis Trabb Pardo²⁶ assign this honor to Alick Glennie’s Autocode, but I, for one, am unable to recognize the sample Autocode program they give as “algebraic,” especially when it is compared to the corresponding Laning and Zierler program.

All of the early “automatic programming” systems were costly to use, since they slowed down the machine by a factor of five or 10. The most common reason for the slowdown was that these systems were spending most of their time in floating-point subroutines. Simulated indexing and other “housekeeping” operations could be done with simple inefficient techniques, since, slow as they were, they took far less time than the floating-point work.

Experience with slow “automatic programming” systems, plus their own experience with the problems of organizing loops and address modification, had convinced programmers that efficient programming was something that could not be automated. Another reason that “automatic programming” was not taken seriously by the computing community was due to the energetic public relations efforts of some visionaries spreading the word that their “automatic programming” systems had almost human abilities to understand the language and the needs of the user, whereas closer inspection of those systems would often reveal a complex, exception-ridden performer of clerical tasks that was both difficult

to use and inefficient. Whatever the reasons, it is difficult to convey to a reader in the late 1970s the strength of the skepticism about “automatic programming” in general and about its ability to produce efficient programs in particular, as it existed in 1954.

In the above discussion of attitudes about “automatic programming” in 1954, I have mentioned only those actual systems of which my colleagues and I were aware at the time. For a comprehensive treatment of early programming systems and languages, I recommend the articles by Knuth and Pardo²⁶ and by Jean Sammet.³⁸

The Economics of Programming

Another factor that influenced the development of Fortran was the economics of programming in 1954. The cost of programmers associated with a computer center was usually at least as great as the cost of the computer itself. (This fact follows from the average salary plus overhead and the number of programmers at each center and from the computer rental figures.) In addition, from one quarter to one half of the computer’s time was spent in debugging. Thus, programming and debugging accounted for as much as three quarters of the cost of operating a computer; and obviously, as computers got cheaper, this situation would get worse.

Fortran is a collection of warts, held together by bits of syntax.

Anon.

This economic factor was one of the prime motivations that led me to propose the Fortran project in a letter to my boss, Cuthbert Hurd, in late 1953 (the exact date is not known, but other facts suggest December 1953 as a likely date). I believe that the economic need for a system like Fortran was one reason why IBM and my successive bosses, Hurd, Charles DeCarlo, and John McPherson, provided for our constantly expanding needs over the next five years without ever asking us to project or justify those needs in a formal budget.

Programming Systems in 1954

It is difficult for a programmer of today to comprehend what “automatic programming” meant to programmers in 1954. To many, it then meant simply providing mnemonic operation codes and symbolic addresses; to others, it meant the simple process of obtaining subroutines from a library and inserting the addresses of operands into each subroutine. Most “automatic programming” systems were assembly programs, subroutine-fixing programs, or, most popularly, interpretive systems to provide floating-point and indexing operations. My friends and I were aware of a number of assembly programs and interpretive systems, some of which have been mentioned above; besides these, there were primarily two other systems of significance: the A-2 compiler^{29,36} and the Laning and Zierler²⁷ algebraic compiler at MIT. As noted above, the A-2 compiler was, at that time, largely a subroutine fixer (its other principal task was to provide for “overlays”); but from the standpoint of its input “programs,” it provided fewer conveniences than most of the then-current interpretive systems mentioned earlier; it

later adopted a “pseudocode” as input that was similar to the input codes of these interpretive systems.

The Laning and Zierler system accepted as input an elegant, but rather simple algebraic language. It permitted single-letter variables (identifiers) that could have a single constant or variable subscript. The repertoire of functions one could use was denoted by “F” with an integer superscript to indicate the “catalog number” of the desired function. Algebraic expressions were compiled into closed subroutines and placed on a magnetic drum for subsequent use. The system was originally designed for the Whirlwind computer when it had 1,024 storage cells, with the result that it caused a slowdown in execution speed by a factor of about 10.¹

The effect of the Laning and Zierler system on the development of Fortran is a question that has been muddled by many misstatements on my part. For many years, I believed that we had gotten the idea for using algebraic notation in Fortran from seeing a demonstration of the Laning and Zierler system at MIT. In preparing a paper⁸ for the International Research Conference on the History of Computing at Los Alamos, New Mexico, 10–15 June 1976, I reviewed the matter with Irving Ziller and obtained a copy of a 1954 letter⁶ (which Dr. Laning kindly sent to me). As a result, the facts of the matter have become clear. The letter in question is one I sent to Dr. Laning asking for a demonstration of his system. It makes clear that we had learned of his work at the Office of Naval Research Symposium on Automatic Programming for Digital Computers, 13–14 May 1954, and that the demonstration took place on 2 June 1954. The letter also makes clear that the Fortran project was well under way when the letter was sent (21 May 1954) and included Harlan Herrick, Robert A. Nelson, and Ziller, as well as myself. Furthermore, an article in the proceedings of that same Office of Naval Research (ONR) symposium by Herrick and myself¹¹ shows clearly that we were already considering input expressions like “ $\sum a_{ij} \cdot b_{jk}$ ” and “ $X + Y$.” We went on to raise the question: “Can a machine translate a sufficiently rich mathematical language into a sufficiently economical program at a sufficiently low cost to make the whole affair feasible?”

These and other remarks in our paper presented at the symposium in May 1954 make it clear that we were already considering algebraic input considerably more sophisticated than that of Laning and Zierler’s system when we first heard of their pioneering work. Thus, although Laning and Zierler had already produced the world’s first algebraic compiler, our basic ideas for Fortran had been developed independently; thus, it is difficult to know what, if any new ideas we got from seeing the demonstration of their system.

(In response to suggestions of the Program Committee, let me try to deal explicitly with the question of what work might have influenced our early ideas for Fortran, although it is mostly a matter of listing work of which we were then unaware. I have already discussed the work of Laning and Zierler and the A-2 compiler. The work of Heinz Rutishauser³⁷ [1952] is discussed later. Like most of the world (except perhaps Rutishauser and Corrado Böhm, the latter of whom was the first to describe a compiler in its own language¹⁵), we were entirely unaware of the work of Konrad Zuse.^{41,42} Zuse’s “Plankalkül,” which he completed in 1945, was, in some ways, a more elegant and advanced programming language than those that appeared 10 and 15 years later. We were also unaware of the work of John Mauchly et al. (“Short Code,” 1950), Arthur Burks (“Intermediate PL,” 1950), Böhm (1951), and Glennie (“Autocode,” 1952), all as discussed in Knuth and Pardo.²⁶ We were aware of but not influenced by the

The History of Fortran I, II, and III

automatic programming efforts that simulated a synthetic computer (e.g., MIT's "Summer Session Computer," SHACO, DUAL, Speedcoding, and the Illiac system), since their languages and systems were so different from those of Fortran. Also, we were not influenced by algebraic systems that were designed after our "Preliminary Report"³⁴ but that began operation before Fortran (e.g., BACAIC,²² IT,³³ and Mathe-Matic⁴). Although PACT I¹³ was not an algebraic compiler, it deserves mention as a significant development; it was designed after the Fortran language but was in operation before Fortran; it did not influence our work.)

Our ONR symposium article¹¹ also makes clear that the Fortran group was already aware that it faced a new kind of problem in automatic programming.

The one central attribute of Fortran is its name.

Martin Greenfield

The viability of most compilers and interpreters prior to Fortran had rested on the fact that most source language operations were not machine operations. Thus, even large inefficiencies in performing both looping/testing operations and computing addresses were masked by most of the operating time being spent in floating-point subroutines. But the advent of the IBM 704 with built-in floating point and indexing radically altered the situation. The 704 presented a double challenge to those who wanted to simplify programming; first, it removed the *raison d'être* of earlier systems by providing in hardware the operations they existed to provide; second, it increased the problem of generating efficient programs by an order of magnitude by speeding up floating-point operations by a factor of 10 and, thereby, leaving inefficiencies nowhere to hide. In view of the widespread skepticism about the possibility of producing efficient programs with an automatic programming system and the fact that inefficiencies could no longer be hidden, we were convinced that the kind of system we had in mind would be widely used only if we could demonstrate that it would produce programs almost as efficient as hand-coded ones and do so on virtually every job.

It was our belief that if Fortran, during its first months, were to translate any reasonable "scientific" source program into an object program only half as fast as its hand-coded counterpart, then acceptance of our system would be in serious danger. This belief caused us to regard the design of the translator as the real challenge, not the simple task of designing the language. Our belief in the simplicity of language design was partly confirmed by the relative ease with which similar languages had been independently developed by Rutishauser,³⁷ Laning and Zierler,²⁷ and ourselves, whereas we were alone in seeking to produce very efficient object programs.

To this day, I believe that our emphasis on object program efficiency rather than on language design was basically correct. I believe that had we failed to produce efficient programs, the widespread use of languages like Fortran would have been seriously delayed. In fact, I believe that we are in a similar, but unrecognized situation today: In spite of all the fuss that has been made over myriad language details, current conventional languages are still very weak programming aids, and far more powerful lan-

guages would be in use today if anyone had found a way to make them run with adequate efficiency. In other words, the next revolution in programming will take place only when *both* of the following requirements have been met:

- a new kind of programming language, far more powerful than those of today, has been developed and
- a technique has been found for executing its programs at not much greater cost than that of today's programs.

Because of our 1954 view that success in producing efficient programs was more important than the design of the Fortran language, I consider the history of the compiler construction and the work of its inventors an integral part of the history of the Fortran language; therefore, a later section deals with that subject.

Early Stages of the Fortran Project

After Hurd approved my proposal to develop a practical automatic programming system for the 704 in December 1953 or January 1954, Ziller was assigned to the project. We started work in one of the many small offices the project was to occupy in the vicinity of IBM headquarters at 590 Madison Avenue in New York; the first of these was in the Jay Thorpe Building on Fifth Avenue. By May 1954, we had been joined by Herrick and then by a new employee who had been hired to do technical typing, Nelson (with Ziller, he soon began designing one of the most sophisticated sections of the compiler; he is now an IBM Fellow). By about May, we had moved to the 19th floor of the annex of 590 Madison Avenue, next to the elevator machinery; the ground floor of this building housed the 701 installation on which customers tested their programs before the arrival of their own machines. It was here that most of the Fortran language was designed, mostly by Herrick, Ziller, and myself, except that most of the input-output language and facilities were designed by Roy Nutt, an employee of United Aircraft Corp. who was soon to become a member of the Fortran project. After we had finished designing most of the language, we heard about Rutishauser's proposals for a similar language.³⁷ It was characteristic of the unscholarly attitude of most programmers then, and of ourselves in particular, that we did not bother to carefully review the sketchy translation of his proposals that we finally obtained, since from their symbolic content, they did not appear to add anything new to our proposed language. Rutishauser's language had a **for** statement and one-dimensional arrays, but no IF, GOTO, or I/O statements. Subscript variables could not be used as ordinary variables, and operator precedence was ignored. His 1952 article described two compilers for this language (for more details, see the article by Knuth and Pardo²⁶).

As far as we were aware, we simply made up the language as we went along. We did not regard language design as a difficult problem, merely a simple prelude to the real problem: designing a compiler that could produce efficient programs. Of course, one of our goals was to design a language that would make it possible for engineers and scientists to write programs themselves for the 704. We also wanted to eliminate a lot of the bookkeeping and detailed, repetitive planning that hand-coding involved. Very early in our work, we had in mind the notions of assignment statements, subscripted variables, and the DO statement (which, I believe, was proposed by Herrick). We felt that these provided a good basis for achieving our goals for the language. Whatever else was needed

emerged as we tried to build a way of programming on these basic ideas.

We certainly had no idea that languages almost identical to the one we were working on would be used for more than one IBM computer, not to mention those of other manufacturers. (After all, there were very few computers around back then.) But we did expect our system to have a big impact, in the sense that it would make programming for the 704 very much faster, cheaper, and more reliable. We also expected that if we were successful in meeting our goals, other groups and manufacturers would follow our example in reducing the cost of programming by providing similar systems with different, but similar languages.³⁴

**Fortran—"the infantile disorder"—is
hopelessly inadequate for whatever
computer application you have in mind
today...too clumsy, too risky,
and too expensive.**

Edsger Dijkstra

By the fall of 1954, we had become the "Programming Research Group," and I had become its "manager." By November of that year, we had produced a paper: "Preliminary Report, Specifications for the IBM Mathematical FORMula TRANslating System, Fortran,"³⁴ dated 10 November 1954. In its introduction, we noted that "systems which have sought to reduce the job of coding and debugging problems have offered the choice of easy coding and slow execution or laborious coding and fast execution." On the basis more of faith than of knowledge, we suggested that programs "will be executed in about the same time that would be required had the problem been laboriously hand coded." In what turned out to be a true statement, we said that "Fortran may apply complex, lengthy techniques in coding a problem which the human coder would have neither the time nor inclination to derive or apply."

The language described in the "Preliminary Report" had:

- variables of one or two characters in length,
- function names of three or more characters,
- recursively defined "expressions,"
- subscripted variables with up to three subscripts,
- "arithmetic formulas" (which turn out to be assignment statements), and
- "DO formulas."

These DO formulas could specify both the first and last statements to be controlled, thus permitting a DO to control a distant sequence of statements, as well as specifying a third statement to which control would pass following the end of the iteration. If only one statement was specified, the "range" of the DO was the sequence of statements following the DO down to the specified statement.

Expressions in "arithmetic formulas" could be "mixed": involving both "fixed-point" (integer) and "floating-point" quantities. The arithmetic used (all integer or all floating point) to evaluate a mixed expression was determined by the type of the variable on the left of the "=" sign. "IF formulas" employed an equality or

inequality sign ("=" or ">" or ">=") between two (restricted) expressions, followed by two statement numbers: one for the "true" case, the other for the "false" case.

A "Relabel formula" was designed to make it easy to rotate, say, the indices of the rows of a matrix so that the same computation would apply, after relabeling, even though a new row had been read in and the next computation was now to take place on a different, rotated set of rows. Thus, for example, if b is a 4×4 matrix, after RELABEL (3, 1), a reference to

- $b(1, j)$ has the same meaning as $b(3, j)$ before relabeling;
- $b(2, j)$ after = $b(4, j)$ before;
- $b(3, j)$ after = $b(1, j)$ before; and
- $b(4, j)$ after = $b(2, j)$ before relabeling.

The input-output statements that were provided included the basic notion of specifying the sequence in which data were to be read in or out, but did not include any "Format" statements.

The report also lists four kinds of "specification sentences":

- 1) "dimension sentences" for giving the dimensions of arrays,
- 2) "equivalence sentences" for assigning the same storage locations to variables,
- 3) "frequency sentences" for indicating estimated relative frequency of branch paths or loops to help the compiler optimize the object program, and
- 4) "relative constant sentences" to indicate subscript variables, which are expected to change their values very infrequently.

Toward the end of the report (pp. 26–27), there is a section "Future Additions to the Fortran System." Its first item is: "a variety of new input-output formulas which would enable the programmer to specify various formats for cards, printing, input tapes and output tapes." It is believed that this item is a result of our early consultations with Nutt. This section goes on to list other proposed facilities to be added:

- complex and double precision arithmetic,
- matrix arithmetic,
- sorting,
- solving simultaneous equations,
- differential equations, and
- linear programming problems.

It also describes function definition capabilities similar to those that later appeared in Fortran II:

- facilities for numerical integration,
- a summation operator, and
- table lookup facilities.

The final section of the report (pp. 28–29) discusses programming techniques to use to help the system produce efficient programs. It discusses how to use parentheses to help the system identify identical subexpressions within an expression and thereby eliminate their duplicate calculation. These parentheses had to be supplied only when a recurring subexpression occurred as part of a term (e.g., if $a * b$ occurred in several places, it would be better to write the term $a * b * c$ as $(a * b) * c$ to avoid duplicate calculation); otherwise the system would identify duplicates without any assistance. It also observes that the system would not produce optimal code for loops constructed without DO statements.

This final section of the report also notes that "no special pro-

The History of Fortran I, II, and III

visions have been included in the Fortran system for locating errors in formulas.” It suggests checking a program “by independently recreating the specifications for a problem from its Fortran formulation [!].” It says nothing about the system catching syntactic errors, but notes that an error-finding program can be written after some experience with errors has been accumulated.

Unfortunately, we were hopelessly optimistic in 1954 about the ease of debugging Fortran programs (thus, we find on page 2 of the report: “Since Fortran should virtually eliminate coding and debugging ... [!]”), and, hence, syntactic error-checking facilities in the first distribution of Fortran I were weak. Better facilities were added not long after distribution, and fairly good syntactic checking was provided in Fortran II.

God is real (unless otherwise declared in an explicit type statement or in an implicit declaration).

B. Graham

The Fortran language described in the *Programmer's Reference Manual*,²⁴ dated 15 October 1956, differed in a few respects from that of the Preliminary Report, but considering our ignorance in 1954 of the problems we would later encounter in producing the compiler, there were:

- remarkably few deletions (the Relabel and Relative Constant statements);
- a few retreats, some fortunate, some not (simplification of DO statements, dropping inequalities from IF statements for lack of a “>” symbol, and prohibiting most “mixed” expressions and subscripted subscripts); and
- the rectification of a few omissions (addition of FORMAT, CONTINUE, computed and assigned GOTO statements, increasing the length of variables to up to six characters, and general improvements of input-output statements).

Since our entire attitude about language design had always been a very casual one, the changes we felt to be desirable during the course of writing the compiler were made equally casually.

We never felt that any of them involved a real sacrifice in convenience or power (with the possible exception of the Relabel statement, whose purpose was to coordinate input-output with computations on arrays, but this was one facility that we felt would have been very difficult to implement). I believe the simplification of the original DO statement resulted from the realizations that:

- 1) it would be hard to describe precisely,
- 2) it was awkward to compile, and
- 3) it provided little power beyond that of the final version.

In our naive unawareness of language design problems—of course, we knew nothing of many issues that were later thought to be important, e.g., block structure, conditional expressions, and type declarations—it seemed to us that once one had the notions of the assignment statement, the subscripted variable, and the DO statement in hand (and these were among our earliest ideas), then the remaining problems of language design were trivial. Their solution was thrust on us either by the need to provide some ma-

chine facility such as reading input or by some programming task that could not be done with existing structures. (For example, skip to the end of a DO loop without skipping the indexing instructions there. This gave rise to the CONTINUE statement.)

One much-criticized design choice in Fortran concerns the use of spaces: Blanks were ignored, even blanks in the middle of an identifier. Nutt reminds me that that choice was partly in recognition of a problem widely known in SHARE (the 704 users' association). There was a common problem with keypunchers not recognizing or properly counting blanks in handwritten data, and this caused many errors. We also regarded ignoring blanks as a device to enable programmers to arrange their programs in a more readable form without altering their meaning or introducing complex rules for formatting statements.

Another debatable design choice was to rule out “mixed”-mode expressions involving both integer and floating-point quantities. Although our Preliminary Report had included such expressions and rules for evaluating them, we felt that if code for type conversion were to be generated, the user should be aware of that, and the best way to ensure that he was aware was to ask him to specify them. I believe we were also doubtful of the usefulness of the rules in our report for evaluating mixed expressions. In any case, the most common sort of “mixture” was allowed: Integer exponents and function arguments were allowed in a floating-point expression.

In late 1954 and early 1955, after completing the Preliminary Report, Herrick, Ziller, and I gave perhaps five or six talks about our plans for Fortran to various groups of IBM customers who had ordered a 704 (the 704 had been announced about May 1954). At these talks, we covered the material in the report and discussed our plans for the compiler (which was to be completed within about six months [!]; this was to remain the interval to completion whenever we were asked, until it actually was completed over two years later, in April 1957). In addition to informing customers about our plans, another purpose of these talks was to assemble a list of their objections and further requirements. In this we were disappointed; our listeners were mostly skeptical; I believe they had heard too many glowing descriptions of what turned out to be clumsy systems to take us seriously. In those days, one was accustomed to finding lots of peculiar, but significant restrictions in a system when it finally arrived that had not been mentioned in its original description. Most of all, our claims that we would produce efficient object programs were disbelieved. Whatever the reasons, we received almost no suggestions or feedback; our listeners had done almost no thinking about the problems we faced and had almost no suggestions or criticisms. Thus, we felt that our trips to Washington, D.C., Albuquerque, Pittsburgh, Los Angeles, and one or two other places were not very helpful.

One trip to give our talk, probably in January 1955, had an excellent payoff. This talk, at United Aircraft Corp., resulted in an agreement between our group and Walter Ramshaw at United Aircraft that Nutt would become a regular part of our effort (although remaining an employee of United Aircraft) to contribute his expertise on input-output and assembly routines. With a few breaks due to his involvement in writing various SHARE programs, he would thenceforth come to New York two or three times a week until early 1957.

It is difficult to assess the influence the early work of the Fortran group had on other projects. Certainly the discussion of Lan-

ing and Zierler's algebraic compiler at the ONR symposium in May 1954 would have been more likely to persuade someone to undertake a similar line of effort than would the brief discussion of the merits of using "a fairly natural mathematical language" that appeared there in the paper by Herrick and myself.¹¹ But it was our impression that our discussions with various groups after that time, their access to the Preliminary Report, and their awareness of the extent and seriousness of our efforts, that these factors either gave the initial stimulus to some other projects or at least caused them to be more active than they might have been otherwise. It was our impression, for example, that the "IT" project³³ at Purdue University and later at Carnegie-Mellon University began shortly after the distribution of our Preliminary Report, as did the "Math-Matic" project⁴ at Sperry Rand.

I don't know what the language of the year 2000 will look like, but I know it will be called Fortran.

Tony Hoare

It is not clear what influence, if any, our Los Angeles talk and earlier contacts with members of their group had on the PACT I effort,¹³ which I believe was already in its formative stages when we got to Los Angeles. It is clear that whatever influence the specifications for Fortran may have had on other projects in 1954–1956, our plans were well-advanced and quite firm by the end of 1954, before we had contact or knowledge of those other projects. Our specifications were not affected by them in any significant way, as far as I am aware, even though some were operating before Fortran (since they were primarily interested in providing an input language rather than an optimization, their task was considerably simpler than ours).

The Construction of the Compiler

The Fortran compiler (or "translator," as we called it then) was begun in early 1955, although a lot of work on various schemes that would be used in it had been done in 1954. Herrick had done a lot of trial programming to test our language, and we had worked out the basic sort of machine programs that we wanted the compiler to generate for various source language phrases. Ziller and I had worked out a basic scheme for translating arithmetic expressions.

But the real work on the compiler got under way in our third location on the fifth floor of 15 East 56th Street in New York. By the middle of February, three separate efforts were under way. The first two of these concerned sections 1 and 2 of the compiler, and the third concerned the input, output, and assembly programs we were going to need (see below). We believed back then that these efforts would produce most of the compiler.

(The entire project was carried out by a loose cooperation between autonomous, separate groups of one, two, or three people; each group was responsible for a "section" of the compiler; each group gradually developed and agreed on its own input and output specifications with the groups for neighboring sections; each group invented and programmed the necessary techniques for doing its assigned job.)

Section 1 was to read the entire source program, compile what instructions it could, and file all the rest of the information from the source program in appropriate tables. Thus, the compiler was "one pass," in the sense that it "saw" the source program only once. Herrick was responsible for creating most of the tables, Peter Sheridan (who had recently joined us) compiled all the arithmetic expressions, and Nutt compiled and/or filed the I/O statements. Herrick, Sheridan, and Nutt got some help later on from R.J. Beeber and H. Stern, but they were the architects of section 1 and wrote most of its code. Sheridan devised and implemented a number of optimizing transformations on expressions³⁹ that sometimes radically altered them (of course, without changing their meaning). Nutt transformed the I/O "lists of quantities" into nests of DO statements, which were then treated by the regular mechanisms of the compiler. The rest of the I/O information he filed for later treatment in section 6, the assembler section. (For further details about how the various sections of the compiler worked, see the 1957 paper by Backus et al.¹²)

Using the information that was filed in section 1, section 2 faced a completely new kind of problem; it was required to analyze the entire structure of the program in order to generate optimal code from DO statements and references to subscripted variables. The simplest way to effect a reference to $A(I, J)$ is to evaluate an expression involving the address of $A(1, 1)$, I , and $K \times J$, where K is the length of a column (when A is stored column-wise). But this calculation, with its multiplication, is much less efficient than the way most hand-coded programs effect a reference to $A(I, J)$, namely, by adding an appropriate constant to the address of the preceding reference to the array A whenever I and J are changing linearly. To employ this far-more-efficient method, section 2 had to determine when the surrounding program was changing I and J linearly.

Thus, one problem was that of distinguishing between, on the one hand, references to an array element that the translator might treat by incrementing the address used for a previous reference and those array references, on the other hand, that would require an address calculation starting from scratch with the current values of the subscripts.

It was decided that it was not practical to track down and identify linear changes in subscripts resulting from assignment statements. Thus, the sole criterion for linear changes, and hence for efficient handling of array references, was to be that the subscripts involved were being controlled by DO statements. Despite this simplifying assumption, the number of cases that section 2 had to analyze in order to produce optimal or near-optimal code was very large. (The number of such cases increased exponentially with the number of subscripts; this was a prime factor in our decision to limit them to three; the fact that the 704 had only three index registers was not a factor.)

It is beyond the scope of this paper to go into the details of the analysis that section 2 carried out. It will suffice to say that it produced code of such efficiency that its output would startle the programmers who studied it. It moved code out of loops where that was possible; it took advantage of the differences between row-wise and column-wise scans; it took note of special cases to optimize even the exits from loops. The degree of optimization performed by section 2 in its treatment of indexing, array references, and loops was not equaled until optimizing compilers began to appear in the middle and late 1960s.

The History of Fortran I, II, and III

The architecture and all the techniques employed in section 2 were invented by Nelson and Ziller. They planned and programmed the entire section. Originally, it was their intention to produce the complete code for their area, including the choice of the index register to be used (the 704 had three index registers). When they started looking at the problem, it rapidly became clear that it was not going to be easy to treat it optimally. At that point, I proposed that they should produce a program for a 704 with an unlimited number of index registers and that later sections would analyze the frequency of execution of various parts of the program (by a Monte Carlo simulation of its execution) and then make index register assignments so as to minimize the transfers of items between the store and the index registers.

Fortran is a language to avoid—unless you want some answers.

Anon.

This proposal gave rise to two new sections of the compiler that we had not anticipated, sections 4 and 5 (section 3 was added to convert the output of sections 1 and 2 to the form required for sections 4, 5, and 6). In the fall of 1955, Lois Mitchell Haibt joined our group to plan and program section 4:

- to analyze the flow of a program produced by sections 1 and 2;
- to divide the program into “basic blocks” (which contained no branching);
- to do a Monte Carlo (statistical) analysis of the expected frequency of execution of basic blocks—by simulating the behavior of the program and keeping counts of the use of each block—using information from DO statements and FREQUENCY statements; and
- to collect information about index register usage (for more details, see the paper by Backus et al.¹² and the one by John Cocke and J.T. Schwartz^{19,p.511}).

Section 5 would then do the actual transformation of the program from one having an unlimited number of index registers to one having only three. Again, the section was entirely planned and programmed by Haibt.

Section 5 was planned and programmed by Sheldon Best, who was loaned to our group by agreement with his employer, Charles W. Adams at the Digital Computer Laboratory at MIT. During his stay with us, Best was a temporary IBM employee. Starting in the early fall of 1955, he designed what turned out to be, along with section 2, one of the most intricate and complex sections of the compiler, one that had perhaps more influence on the methods used in later compilers than any other part of the Fortran compiler. (For a discussion of his techniques, see the paper by Cocke and Schwartz.^{19,pp.510-515}) It is impossible to describe his register allocation method here; it suffices to say that it was to become the basis for much subsequent work and produced code very difficult to improve.

Although I believe that no provably optimal register allocation algorithm is known for general programs with loops etc., empirically, Best’s 1955–1956 procedure appeared to be optimal. For straight-line code, Best’s replacement policy was the same as that

used in L.A. Belady’s MIN algorithm, which Belady proved to be optimal.¹⁴ Although Best did not publish a formal proof, he had convincing arguments for the optimality of his policy in 1955.

Late in 1955, it was recognized that yet another section, section 3, was needed. This section merged the outputs of the preceding sections into a single uniform 704 program that could refer to any number of index registers. It was planned and programmed by Richard Goldberg, a mathematician who joined us in November 1955. Also, late in 1956, after Best had returned to MIT and during the debugging of the system, section 5 was taken over by Goldberg and David Sayre (see below), who diagrammed it carefully and took charge of its final debugging.

The final section of the compiler, section 6, assembled the final program into a relocatable binary program (it could also produce a symbolic program in SAP, the SHARE Assembly Program for the 704). It produced a storage map of the program and the data that was a compact summary of the Fortran output. Of course, it also obtained the necessary library programs for inclusion in the object program, including those required to interpret FORMAT statements and perform input-output operations. Taking advantage of the special features of the programs it assembled, this assembler was about 10 times faster than SAP. It was designed and programmed by Nutt, who also wrote all the I/O programs and the relocating binary loader for loading object programs.

By the summer of 1956, large parts of the system were working. Sections 1, 2, and 3 could produce workable code, provided no more than three index registers were needed. A number of test programs were compiled and run at this time. Nutt took part of the system to United Aircraft (sections 1, 2, and 3 and the part of section 6 that produced SAP output). This part of the system was productive there from the summer of 1956 until the complete system was available in early 1957.

From late spring of 1956 to early 1957, the pace of debugging was intense; often we would rent rooms in the Langdon Hotel (which disappeared long ago) on 56th Street, sleep there a little during the day, and then stay up all night to get as much use of the computer (in the headquarters annex on 57th Street) as possible.

It was an exciting period; when we later began to get fragments of compiled programs out of the system, we were often astonished at the surprising transformations in the indexing operations and in the arrangement of the computation that the compiler made, changes that made the object program efficient but that we would not have thought to make as programmers ourselves (even though, of course, Nelson or Ziller could figure out how the indexing worked, Sheridan could explain how an expression had been optimized beyond recognition, and Goldberg or Sayre could tell us how section 5 had generated additional indexing operations). Transfers of control appeared that corresponded to no source statement, expressions were radically rearranged, and the same DO statement might produce no instructions in the object program in one context, and in another it would produce many instructions in different places in the program.

By the summer of 1956, what appeared to be the imminent completion of the project started us worrying (for perhaps the first time) about documentation. Sayre, a crystallographer who had joined us in the spring (he had earlier consulted with Best on the design of section 5 and had later begun serving as second-in-command of what was now the “Programming Research Department”), took up the task of writing the *Programmer’s Reference*

Manual.²⁴ It appeared in a glossy cover, handsomely printed, with the date 15 October 1956. It stood for some time as a unique example of a manual for a programming language (perhaps it still does): It had wide margins, yet was only 51 pages long. Its description of the Fortran language, exclusive of input-output statements, was 21 pages; the I/O description occupied another 11 pages; the rest of it was examples and details about arithmetic, tables, etc. It gave an elegant recursive definition of expressions (as given by Sheridan) and concise, clear descriptions, with examples, of each statement type, of which there were 32, mostly machine-dependent items like SENSE LIGHT, IF DIVIDE CHECK, PUNCH, READ DRUM, and so on.

One feature of Fortran I is missing from the *Programmer's Reference Manual*, not from an oversight of Sayre's, but because it was added to the system after the manual was distributed. This feature was the ability to define a function by a "function statement." These statements had to precede the rest of the program. They looked like assignment statements, with the defined function and dummy arguments on the left and an expression involving those arguments on the right. They are described in the addenda to the *Programmer's Reference Manual*,² which we sent on 8 February 1957 to John Greenstadt, who was in charge of IBM's facility for distributing information to SHARE. They are also described in all subsequent material on Fortran I.

**In the good old days, physicists
repeated each other's experiments, just
to be sure. Today, they stick to Fortran
so they can share each others'
programs, bugs included.**

Edsger Dijkstra

The next documentation task we set ourselves was to write a paper describing the Fortran language and the translator program. The result was a paper entitled, "The Fortran Automatic Coding System,"¹² which we presented at the Western Joint Computer Conference in Los Angeles in February 1957. I have mentioned all of the 13 authors of that paper in the preceding narrative except one: Robert A. Hughes. He was employed by the Livermore Radiation Laboratory; by arrangement with Sidney Fernbach, he visited us for two or three months in the summer of 1956 to help us document the system.

At about the time of the Western Joint Computer Conference, we spent some time in Los Angeles still frantically debugging the system. North American Aviation gave us time at night on their 704 to help us in our mad rush to distribute the system. Up to this point, there had been relatively little interest from 704 installations (with the exception of Ramshaw's United Aircraft shop, Harry Cantrell's GE installation in Schenectady, New York, and Fernbach's Livermore operation), but now that the full system was beginning to generate object programs, interest picked up in a number of places.

Sometime in early April 1957, we felt the system was sufficiently bug-free to distribute to all 704 installations. Sayre and Grace E. Mitchell (see below) started to punch out the binary decks of the system, each of about 2,000 cards, with the intention

to make 30 or 40 decks for distribution. This process was so error-prone that they had to give up after spending an entire night in producing only one or two decks. Apparently one of those decks was sent, without any identification or directions, to the Westinghouse-Bettis installation, where a puzzled group headed by Herbert S. Bright, suspecting that it might be the long-awaited Fortran deck, proceeded, entirely by guesswork, to get it to compile a test program—after a diagnostic printout noted that a comma was missing in a specific statement! This program then printed 28 pages of correct results—with a few FORMAT errors. The date: 20 April 1957. Bright wrote an amusing account of this incident in *Computers and Automation*.¹⁷ [Editor's note: This story is also told in *Annals*, vol. 1, no. 1, pp. 72–74.]

After failing to produce binary decks, Sayre devised and programmed the simple editor and loader that made it possible to distribute and update the system from magnetic tapes; this arrangement served as the mechanism for creating new system tapes from a master tape and the binary correction cards that our group would generate in large numbers during the long field debugging and maintenance period that followed distribution.

With the distribution of the system tapes went a *Preliminary Operator's Manual*³¹ dated 8 April 1957. It describes how to use the tape editor and how to maintain the library of functions. Five pages of such general instructions are followed by 32 pages of error stops, many of these say, "source program error, get off machine, correct formula in question and restart problem" and then, for example, (stop 3,624) "non-zero level reduction due to insufficient or redundant parentheses in arithmetic or IF-type formula." Shortly after the distribution of the system, we distributed—one copy per installation—what was fondly known as the "Tome," the complete symbolic listing of the entire compiler plus other system and diagnostic information, an 11 × 15 inch volume about four or five inches thick.

The proprietors of the six sections were kept busy tracking down bugs elicited by our customers' use of Fortran until the late summer of 1957. Hal Stern served as the coordinator of the field debugging and maintenance effort; he received a stream of telegrams, mail, and phone calls from all over the country and distributed the incoming problems to the appropriate members of our group to track down the errors and generate correction cards, which he then distributed to every installation.

In the spring of 1957, Mitchell joined our group to write the *Programmer's Primer* for Fortran. The primer was divided into three sections; each described successively larger subsets of the language, accompanied by many example programs. The first edition of the primer was issued in the late fall or winter of 1957; a slightly revised edition appeared in March 1958. Mitchell planned and wrote the 64-page primer with some consultation with the rest of the group; she later programmed most of the extensive changes in the system that resulted in Fortran II (see below).

The primer had an important influence on the subsequent growth in the use of the system. I believe it was the only available simplified instruction manual (other than reference manuals) until the later appearance of books such as those by Daniel McCracken²⁸ (1961), Elliot Organick³² (1963), and many others.

A report on Fortran usage in November 1958⁷ says that "a survey in April [1958] of twenty-six 704 installations indicates that over half of them use Fortran [I] for more than half of their problems. Many use it for 80% or more of their work ... and almost all

The History of Fortran I, II, and III

use it for some of their work.” By the fall of 1958, there were some 60 installations with about 66 704s and “more than half the machine instructions for these machines are being produced by Fortran. SHARE recently designated Fortran as the second official medium for transmittal of programs [SAP was the first].”

Fortran II

During the field debugging period, some shortcomings of the system design, which we had been aware of earlier but had no time to correct, were constantly coming to our attention. In the early fall of 1957, we started to plan ways of correcting these shortcomings; a document dated 25 September 1957³⁵ characterizes them as:

- 1) the need for better diagnostics (clearer comments about the nature of source program errors) and
- 2) the need for subroutine definition capabilities.

(Although one Fortran I diagnostic would pinpoint, in a print-out, a missing comma in a particular statement, others could be very cryptic.) This document is titled, “Proposed Specifications for Fortran II for the 704”; it sketches a more general diagnostic system and describes the new “subroutine definition” and END statements, plus some others that were not implemented. It describes how symbolic information is retained in the relocatable binary form of a subroutine so that the “binary symbolic subroutine [BSS] loader” can implement references to separately compiled subroutines. It describes new prologues for these subroutines and points out that mixtures of Fortran-coded and assembly-coded relocatable binary programs could be loaded and run together. It does not discuss the FUNCTION statement that was also available in Fortran II. Fortran II was designed mostly by Nelson, Ziller, and myself. Mitchell programmed the majority of new code for Fortran II (with the most unusual feature that she delivered it ahead of schedule). She was aided in this by Bernyce Brady and LeRoy May. Sheridan planned and made the necessary changes in his part of section 1; Nutt did the same for section 6. Fortran II was distributed in the spring of 1958.

Fortran III

While Fortran II was being developed, Ziller was designing an even more advanced system that he called Fortran III. It allowed one to write intermixed symbolic instructions and Fortran statements. The symbolic (704) instructions could have Fortran variables (with or without subscripts) as “addresses.” In addition to this machine-dependent feature (which assured the demise of Fortran III along with that of the 704), it contained early versions of a number of improvements that were later to appear in Fortran IV. It had “Boolean” expressions, function and subroutine names could be passed as arguments, and it had facilities for handling alphanumeric data, including a new FORMAT code “A” similar to codes “I” and “E.” This system was planned and programmed by Ziller with some help from Nelson and Nutt. Ziller maintained it and made it available to about 20 (mostly IBM) installations. It was never distributed generally. It was accompanied by a brief descriptive document.³ It became available on this limited scale in the winter of 1958–1959 and was in operation until the early 1960s, in part on the 709 using the compatibility feature (which made the 709 order code the same as that of the 704).

Fortran After 1958, Comments

By the end of 1958 or early 1959, the Fortran group (the Programming Research Department), while still helping with an occasional debugging problem with Fortran II, was primarily occupied with other research. Another IBM department had long since taken responsibility for the Fortran system and was revising it in the course of producing a translator for the 709 that used the same procedures as the 704 Fortran II translator. Since my friends and I no longer had responsibility for Fortran and were busy thinking about other things by the end of 1958, that seems like a good point to break off this account. There remain only a few comments to be made about the subsequent development of Fortran.

The most obvious defect in Fortran II for many of its users was the time spent in compiling. Even though the facilities of Fortran II permitted separate compilation of subroutines and, hence, eliminated the need to recompile an entire program at each step in debugging it, nevertheless, compile times were long and, during debugging, the considerable time spent in optimizing was wasted. I repeatedly suggested to those who were in charge of Fortran that they should now develop a fast compiler and/or interpreter without any optimizing at all for use during debugging and for short-run jobs. Unfortunately, the developers of Fortran IV thought they could have the best of both worlds in a single compiler—one that was both fast and produced optimized code. I was unsuccessful in convincing them that two compilers would have been far better than the compromise that became the original Fortran IV compiler. The latter was not nearly as fast as later compilers like WATFOR²⁰ and did not produce as good code as Fortran II. (For more discussion of later developments with Fortran, see the 1964 paper by Backus and W.P. Heising.¹⁰)

My own opinion as to the effect of Fortran on later languages and the collective impact of such languages on programming generally is not a popular opinion. That viewpoint is the subject of a long paper.⁹ I now regard all conventional languages (e.g., the Fortrans, the Algols, and their successors and derivatives) as increasingly complex elaborations of the style of programming dictated by the von Neumann computer. These “von Neumann languages” create enormous, unnecessary intellectual roadblocks in thinking about programs and in creating the higher-level combining forms required in a powerful programming methodology. The von Neumann languages constantly keep our noses pressed in the dirt of address computation and the separate computation of single words, whereas we should be focusing on the form and content of the overall result we are trying to produce. We have come to regard the DO, FOR, and WHILE statements and the like as powerful tools, whereas they are, in fact, weak palliatives that are necessary to make the primitive von Neumann style of programming viable at all.

By splitting programming into a world of expressions, on the one hand, and a world of statements, on the other, von Neumann languages prevent the effective use of higher-level combining forms; the lack of the latter makes the definitional capabilities of von Neumann languages so weak that most of their important features cannot be defined, starting with a small, elegant framework—but must be built into the framework of the language at the outset. The gargantuan size of recent von Neumann languages is eloquent proof of their inability to define new constructs: No one would build in so many complex features if they could be defined and would fit into the existing framework later.

The world of expressions has some elegant and useful mathematical properties, whereas the world of statements is a disorderly one without useful mathematical properties. Structured programming can be viewed as a modest effort to introduce a small amount of order into the chaotic world of statements. The work of C.A.R. Hoare²³ (1969), Edsger Dijkstra²¹ (1976), and others to axiomatize the properties of the statement world can be viewed as a valiant and effective effort to be precise about those properties, ungainly as they may be.

This is not the place for me to elaborate any further my views about von Neumann languages. My point is this: While it was perhaps natural and inevitable that languages like Fortran and its successors should have developed out of the concept of the von Neumann computer as they did, the fact that such languages have dominated our thinking for 20 years is unfortunate. It is unfortunate because their long-standing familiarity will make it hard for us to understand and adopt new programming styles that one day will offer far greater intellectual and computational power.

Acknowledgments

My greatest debt in connection with this paper is to my old friends and colleagues whose creativity, hard work, and invention made Fortran possible. It is a pleasure to acknowledge my gratitude to them for their contributions to the project, for making our work together so long ago such a congenial and memorable experience, and, more recently, for providing me with a great amount of information and helpful material in preparing this paper and for their careful reviews of an earlier draft. For all this, I thank all those who were associated with the Fortran project but who are too numerous to list here. In particular, I want to thank those who were the principal movers in making Fortran a reality: Sheldon Best, Richard Goldberg, Lois Haiht, Harlan Herrick, Grace Mitchell, Robert Nelson, Roy Nutt, David Sayre, Peter Sheridan, and Irving Ziller. I also wish to thank Bernard Galler, J.A.N. Lee, and Henry Tropp for their amiable, extensive, and invaluable suggestions for improving the first draft of this paper. I am grateful, too, for all the work of the Program Committee in preparing helpful questions that suggested a number of topics in the paper.

References

[Editor's note: Most of the items listed below have dates in the 1950s, thus many that appeared in the open literature will be obtainable only in the largest and oldest collections. The items with an asterisk were either not published or were of such a nature as to make their availability even less likely than that of the other items.]

[1] Adams, Charles W., and Laning, J.H., Jr., "The MIT Systems of Automatic Coding: Comprehensive, Summer Session, and Algebraic," *Proc. Symp. Automatic Programming for Digital Computers*. Washington, D.C.: Office of Naval Research, May 1954.

[2] **Addenda to the Fortran Programmer's Reference Manual*, 8 February 1957. (Transmitted to Dr. John Greenstadt, Special Programs Group, Applied Science Division, IBM, for distribution to SHARE members, by letter from John W. Backus, Programming Research Dept. IBM. 5 pages.)

[3] *Additions to Fortran II, *Description of Source Language Additions to the Fortran II System*. New York: Programming Research, IBM Corp., 1958. (Distributed to users of Fortran III. 12 pages.)

[4] *Ash, R., Broadwin, E., Della Valle, V., Katz, C., Green, M., Jenny, A., and Yu, L., *Preliminary Manual for MATH-MATIC and ARITH-MATIC Systems (for Algebraic Translations and Compilation for UNIVAC I and II)*. Philadelphia: Remington Rand Univac, 1957.

[5] Backus, J.W., "The IBM 701 Speedcoding System," *J. ACM*, vol. 1, no. 1, pp. 4-6, Jan. 1954.

[6] * Backus, John, Letter to J.H. Laning, Jr., 21 May 1954.

[7] Backus, J.W., "Automatic Programming: Properties and Performance of FORTRAN Systems I and II," *Proc. Symp. Mechanisation of Thought Processes*. Teddington, Middlesex, England, National Physical Laboratory, Nov. 1958.

[8] Backus, John, "Programming in America in the Nineteen Fifties—Some Personal Impressions," N. Metropolis, J. Howlett, and Gian-Carlo Rota, eds., *A History of Computing in the Twentieth Century (Proc. Int'l Conf. History of Computing)*, Los Alamos, N.M., 1978). New York: Academic Press, in press.

[9] Backus, John, "Can Programming Be Liberated From the von Neumann Style? A Functional Style and Its Algebra of Programs," *Comm. ACM*, vol. 21, no. 8, pp. 613-641, Aug. 1978.

[10] Backus, J.W., and Heising, W.P., "Fortran," *IEEE Trans. Electronic Computers*, vol. 13, no. 4, pp. 382-385, Aug. 1964.

[11] Backus, John W., and Herrick, Harlan, "IBM 701 Speedcoding and Other Automatic Programming Systems," *Proc. Symp. Automatic Programming for Digital Computers*. Washington, D.C.: Office of Naval Research, May 1954.

[12] Backus, J.W., Beeber, R.J., Best, S., Goldberg, R., Haiht, L.M., Herrick, H.L., Nelson, R.A., Sayre, D., Sheridan, P.B., Stern, H., Ziller, I., Hughes, R.A., and Nutt, R., "The Fortran Automatic Coding System," *Proc. Western Joint Computer Conf.*, Los Angeles, Feb. 1957.

[13] Baker, Charles L., "The PACT I Coding System for the IBM Type 701," *J. ACM*, vol. 3, no. 4, pp. 272-278, Oct. 1956.

[14] Belady, L.A., *Measurements on Programs: One Level Store Simulation*. Yorktown Heights, N.Y.: IBM Thomas J. Watson Research Center, Tech. Rep. RC 1420, 15 June 1965.

[15] Böhm, Corrado, "Calculatrices digitales: Du déchiffrement de formules logico-mathématiques par la machine même dans la conception du programme," *Ann. di Mat. Pura ed Applicata*, vol. 37, no. 4, pp. 175-217, 1954.

[16] Bouricius, Willard G., "Operating Experience With the Los Alamos 701," *Proc. Eastern Joint Computer Conf.*, Washington, D.C., Dec. 1953.

[17] Bright, Herbert S., "Fortran Comes to Westinghouse-Bettis," *Computers and Automation*, Nov. 1971. Also in *Annals of the History of Computing*, July 1979.

[18] Brown, J.H. and Carr, John W., III, "Automatic Programming and Its Development on MIDAC," *Proc. Symp. Automatic Programming for Digital Computers*. Washington, D.C.: Office of Naval Research, May 1954.

[19] Cocks, John, and Schwartz, J.T., *Programming Languages and Their Compilers*. New York: New York Univ., Courant Institute of Mathematical Sciences, Apr. 1970.

[20] Cress, Paul, Dirksen, Paul, and Graham, J. Wesley, *Fortran IV With WATFOR and WATFIV*. Englewood Cliffs, N.J.: Prentice Hall, 1970.

[21] Dijkstra, Edsger W., *A Discipline of Programming*. Englewood Cliffs, N.J.: Prentice Hall, 1976.

[22] Grems, Mandalay, and Porter, R.E., "A Truly Automatic Programming System," *Proc. Western Joint Computer Conf.*, 1956.

[23] Hoare, C.A.R., "An Axiomatic Basis for Computer Programming," *Comm. ACM*, vol. 12, pp. 576-580, 583, Oct. 1969.

[24] *IBM, *Programmer's Reference Manual, the Fortran Automatic Coding System for the IBM 704 EDPM*. New York: IBM Corp., 15 Oct. 1956 (Applied Science Division and Programming Research Dept., Working Committee: J.W. Backus, R.J. Beeber, S. Best, R. Goldberg, H.L. Herrick, R.A. Hughes [Univ. of Calif. Radiation Lab., Livermore, Calif.], L.B. Mitchell, R.A. Nelson, R. Nutt [United Aircraft Corp., East Hartford, Conn.] D. Sayre, P.B. Sheridan, H. Stern, I. Ziller).


[25] * IBM, *Programmer's Primer for Fortran Automatic Coding System for the IBM 704*. New York: IBM Corp., form no. 32-0306, 1957.

[26] Knuth, Donald E., and Pardo, Luis Trabb, "Early Development of Programming Languages," *Encyclopedia of Computer Science and Technology*, vol. 7, p. 419. New York: Marcel Dekker, 1977.

[27] *Laning, J.H., and Zierler, N., *A Program for Translation of Mathematical Equations for Whirlwind I*. Cambridge, Mass.: MIT Instrumentation Lab., Engineering Memorandum E-364, Jan. 1954.

The History of Fortran I, II, and III

- [28] McCracken, Daniel D., *A Guide to Fortran Programming*. New York: John Wiley, 1961.
- [29] Moser, Nora B., "Compiler Method of Automatic Programming," *Proc. Symp. Automatic Programming for Digital Computers*. Washington, D.C.: Office of Naval Research, May 1954.
- [30] Muller, David E., "Interpretive Routines in the ILLIAC Library," *Proc. Symp. Automatic Programming for Digital Computers*. Washington, D.C.: Office of Naval Research, May 1954.
- [31] **Preliminary Operator's Manual for the Fortran Automatic Coding System for the IBM 704 EDPM*. New York: IBM Corp., Programming Research Dept., 8 Apr. 1957.
- [32] Organick, Elliot I., *A Fortran Primer*. Reading, Mass.: Addison-Wesley, 1963.
- [33] *Perlis, A.J., Smith, J.W., and Van Zoeren, H.R., *Internal Translator (IT): A Compiler for the 650*. Pittsburgh, Pa.: Carnegie Institute of Technology, Mar. 1957.
- [34] * "Preliminary Report," *Specifications for the IBM Mathematical FORMula TRANslating System, Fortran*. New York: IBM Corp., 10 Nov. 1954 (report by Programming Research Group, Applied Science Division, IBM. Distributed to prospective 704 customers and other interested parties. 29 pp.).
- [35] * "Proposed Specifications," *Proposed Specifications for Fortran II for the 704*, 25 Sept. 1957. (Unpublished memorandum, Programming Research Dept., IBM.)
- [36] *Remington Rand, Inc., *The A-2 Compiler System Operations Manual*. Prepared by Richard K. Ridgway and Margaret H. Harper under the direction of Grace M. Hopper, 15 Nov. 1953.
- [37] Rutishauser, Heinz, "Automatische Rechenplanfertigung bei programmgesteuerten Rechenmaschinen," *Mitteilungen aus dem Inst. für angew. Math. an der E. T. H. Zunch.*, nr. 3. Basel: Birkhauser, 1952.
- [38] Sammet, Jean E., *Programming Languages: History and Fundamentals*. Englewood Cliffs, N.J.: Prentice Hall, 1969.
- [39] Sheridan, Peter B., "The Arithmetic Translator-Compiler of the IBM Fortran Automatic Coding System," *Comm. ACM*, vol. 2, no. 2, pp. 9-21, Feb. 1959.
- [40] *Schlesinger, S.I., *Dual Coding System*. Los Alamos, N.M.: Los Alamos Science Lab., Rep. LA1573, July 1953.
- [41] Zuse, K., "Über den Plankalkül," *Electron Rechenanl.*, vol. 1, pp. 68-71, 1959.
- [42] Zuse, K., "Der Plankalkül," *Berichte der Gesellschaft für Mathematik und Datenverarbeitung*, vol. 63, part 3. Bonn, 1972. (Manuscript prepared in 1945.)



John Backus led the development of Fortran and participated in designing the IBM 704 and Algol, he proposed "BNF" for syntax description. He received the National Medal of Science in 1975, the Turing Award (ACM) in 1977, the IEEE Computer Society Pioneer Award in 1980, and the Charles Stark Draper Award from the National Academy of Engineering in 1993. He is a member of the National Academy of Sciences and the National Academy of Engineering. He is currently retired from his job at the IBM Research Laboratory, San Jose, California, but continues to consult with IBM from time to time.

The author can be contacted at
91 St. Germain Ave.
San Francisco, CA 94114