

# NUMERICAL RECIPES in Fortran 90

*The Art of Parallel Scientific Computing*

William H. Press

Saul A. Teukolsky

William T. Vetterling

Brian P. Flannery

# Numerical Recipes

## in Fortran 90

The Art of *Parallel* Scientific Computing  
Second Edition

Volume 2 of  
Fortran Numerical Recipes

**William H. Press**

*Harvard-Smithsonian Center for Astrophysics*

**Saul A. Teukolsky**

*Department of Physics, Cornell University*

**William T. Vetterling**

*Polaroid Corporation*

**Brian P. Flannery**

*EXXON Research and Engineering Company*

*Foreword by*

**Michael Metcalf**

*CERN, Geneva, Switzerland*



**CAMBRIDGE**  
UNIVERSITY PRESS

Published by the Press Syndicate of the University of Cambridge  
The Pitt Building, Trumpington Street, Cambridge CB2 1RP  
40 West 20th Street, New York, NY 10011-4211, USA  
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

Copyright © Cambridge University Press 1986, 1996,  
except for all computer programs and procedures, which are  
Copyright © Numerical Recipes Software 1986, 1996,  
and except for Appendix C1, which is placed into the public domain.  
All Rights Reserved.

Numerical Recipes in Fortran 90: The Art of Parallel Scientific Computing,  
Volume 2 of Fortran Numerical Recipes, Second Edition, first published 1996.  
The code in this volume is corrected to software version 2.06

Printed in the United States of America  
Typeset in  $\text{\TeX}$

Without an additional license to use the contained software, this book is intended as a text and reference book, for reading purposes only. A free license for limited use of the software by the individual owner of a copy of this book who personally types one or more routines into a single computer is granted under terms described on p. xviii. See the section "License Information" (pp. xvii–xx) for information on obtaining more general licenses at low cost.

Machine-readable media containing the software in this book, with included licenses for use on a single screen, are available from Cambridge University Press. See the order form at the back of the book, email to "orders@cup.org" (North America) or "trade@cup.cam.ac.uk" (rest of world), or write to Cambridge University Press, 110 Midland Avenue, Port Chester, NY 10573 (USA), for further information.

The software may also be downloaded, with immediate purchase of a license also possible, from the Numerical Recipes Software Web site (<http://www.nr.com>). Unlicensed transfer of Numerical Recipes programs to any other format, or to any computer except one that is specifically licensed, is strictly prohibited. Technical questions, corrections, and requests for information should be addressed to Numerical Recipes Software, P.O. Box 243, Cambridge, MA 02238 (USA), email "info@nr.com", or fax 617 863-1739.

*Library of Congress Cataloging-in-Publication Data*

Numerical recipes in Fortran 90 : the art of parallel scientific computing / William H. Press

... [et al.]. -- 2nd ed.

p. cm.

Includes bibliographical references and index.

ISBN 0-521-57439-0 (hardcover)

1. FORTRAN 90 (Computer program language) 2. Parallel programming (Computer science) 3. Numerical analysis--Data processing.

I. Press, William H.  
QA76.73.F25N85 1996  
519.4'0285'52--dc20

96-5567  
CIP

A catalog record for this book is available from the British Library.

ISBN 0 521 57439 0 Volume 2 (this book)  
ISBN 0 521 43064 X Volume 1  
ISBN 0 521 43721 0 Example book in FORTRAN  
ISBN 0 521 57440 4 FORTRAN diskette (IBM 3.5")  
ISBN 0 521 57608 3 CDROM (IBM PC/Macintosh)  
ISBN 0 521 57607 5 CDROM (UNIX)

# Preface to Volume 2

Fortran 90 is not just the long-awaited updating of the Fortran language to modern computing practices. It is also the vanguard of a much larger revolution in computing, that of multiprocessor computers and widespread parallel programming. Parallel computing has been a feature of the largest supercomputers for quite some time. Now, however, it is rapidly moving towards the desktop.

As we watched the gestation and birth of Fortran 90 by its governing “X3J3 Committee” (a process interestingly described by a leading committee member, Michael Metcalf, in the Foreword that follows), it became clear to us that the right moment for moving Numerical Recipes from Fortran 77 to Fortran 90 was sooner, rather than later.

Fortran 90 compilers are now widely available. Microsoft’s Fortran PowerStation for Windows 95 brings that firm’s undeniable marketing force to PC desktop; we have tested this compiler thoroughly on our code and found it excellent in compatibility and performance. In the UNIX world, we have similarly tested, and had generally fine experiences with, DEC’s Fortran 90 for Alpha AXP and IBM’s xlf for RS/6000 and similar machines. NAG’s Fortran 90 compiler also brings excellent Fortran 90 compatibility to a variety of UNIX platforms. There are no doubt other excellent compilers, both available and on the way. Fortran 90 is completely backwards compatible with Fortran 77, by the way, so you don’t have to throw away your legacy code, or keep an old compiler around.

There have been previous special versions of Fortran for parallel supercomputers, but always specific to a particular hardware. Fortran 90, by contrast, is designed to provide a general, architecture-independent framework for parallel computation. Equally importantly, it is an international standard, agreed upon by a large group of computer hardware and software manufacturers and international standards bodies.

With the Fortran 90 language as a tool, we want this volume to be your complete guide for learning how to “think parallel.” The language itself is very general in this regard, and applicable to many present and future computers, or even to other parallel computing languages as they come along. Our treatment emphasizes general principles, but we are also not shy about pointing out parallelization “tricks” that have frequent applicability. These are not only discussed in this volume’s principal text chapters (Chapters 21–23), but are also sprinkled throughout the chapters of Fortran 90 code, called out by a special “parallel hint” logo (left, above). Also scattered throughout the code chapters are specific “Fortran 90 tips,” with their own distinct graphic call-out (left). After you read the text chapters, you might want simply to browse among these hints and tips.

A special note to C programmers: Right now, there is no effort at producing a parallel version of C that is comparable to Fortran 90 in maturity, acceptance, and stability. We think, therefore, that C programmers will be well served by using this volume for an educational excursion into Fortran 90, its parallel programming constructions, and the numerical algorithms that capitalize on them. C and C++ programming have not been far from our minds as we have written this volume, and we think that you will find that time spent in absorbing its principal lessons (in Chapters 21–23) will be amply repaid in the future, as C and C++ eventually develop standard parallel extensions.

A final word of truth in packaging: **Don't buy this volume unless you also buy (or already have) Volume 1** (now retitled *Numerical Recipes in Fortran 77*). Volume 2 does not repeat any of the discussion of what individual programs actually do, or of the mathematical methods they utilize, or how to use them. While our Fortran 90 code is thoroughly commented, and includes a header comment for each routine that describes its input and output quantities, these comments are *not* supposed to be a complete description of the programs; the complete descriptions are in Volume 1, which we reference frequently. But here's a money-saving hint to our previous readers: If you already own a Second Edition version whose title is *Numerical Recipes in FORTRAN* (which doesn't indicate either "Volume 1" or "Volume 2" on its title page) then take a marking pen and write in the words "Volume 1." There! (Differences between the previous reprintings and the newest reprinting, the one labeled "Volume 1," are minor.)

## Acknowledgments

We continue to be in the debt of many colleagues who give us the benefit of their numerical and computational experience. Many, though not all, of these are listed by name in the preface to the second edition, in Volume 1. To that list we must now certainly add George Marsaglia, whose ideas have greatly influenced our new discussion of random numbers in this volume (Chapter B7).

With this volume, we must acknowledge our additional gratitude and debt to a number of people who generously provided advice, expertise, and time (a great deal of time, in some cases) in the areas of parallel programming and Fortran 90. The original inspiration for this volume came from Mike Metcalf, whose clear lectures on Fortran 90 (in this case, overlooking the beautiful Adriatic at Trieste) convinced us that Fortran 90 could serve as the vehicle for a book with the larger scope of parallel programming generally, and whose continuing advice throughout the project has been indispensable. Gyan Bhanot also played a vital early role in the development of this book; his first translations of our Fortran 77 programs taught us a lot about parallel programming. We are also grateful to Greg Lindhorst, Charles Van Loan, Amos Yahil, Keith Kimball, Malcolm Cohen, Barry Caplin, Loren Meissner, Mitsu Sakamoto, and George Schnurer for helpful correspondence and/or discussion of Fortran 90's subtler aspects.

We once again express in the strongest terms our gratitude to programming consultant Seth Finkelstein, whose contribution to both the coding and the thorough testing of all the routines in this book (against multiple compilers and in sometimes-buggy, and always challenging, early versions) cannot be overstated.

WHP and SAT acknowledge the continued support of the U.S. National Science Foundation for their research on computational methods.

February 1996

William H. Press  
Saul A. Teukolsky  
William T. Vetterling  
Brian P. Flannery

# Foreword

by Michael Metcalf

Sipping coffee on a sunbaked terrace can be surprisingly productive. One of the *Numerical Recipes* authors and I were each lecturing at the International Center for Theoretical Physics in Trieste, Italy, he on numerical analysis and I on Fortran 90. The numerical analysis community had made important contributions to the development of the new Fortran standard, and so, unsurprisingly, it became quickly apparent that the algorithms for which *Numerical Recipes* had become renowned could, to great advantage, be recast in a new mold. These algorithms had, hitherto, been expressed in serial form, first in Fortran 77 and then in C, Pascal, and Basic. Now, nested iterations could be replaced by array operations and assignments, and the other features of a rich array language could be exploited. Thus was the idea of a "Numerical Recipes in Fortran 90" first conceived and, after three years' gestation, it is a delight to assist at the birth.

But what *is* Fortran 90? How did it begin, what shaped it, and how, after nearly foundering, did its driving forces finally steer it to a successful conclusion?

## ***The Birth of a Standard***

Back in 1966, the version of Fortran now known as Fortran 66 was the first language ever to be standardized, by the predecessor of the present American National Standards Institute (ANSI). It was an all-American affair. Fortran had first been developed by John Backus of IBM in New York, and it was the dominant scientific programming language in North America. Many Europeans preferred Algol (in which Backus had also had a hand). Eventually, however, the mathematicians who favored Algol for its precisely expressible syntax began to defer to the scientists and engineers who appreciated Fortran's pragmatic, even natural, style. In 1978, the upgraded Fortran 77 was standardized by the ANSI technical committee, X3J3, and subsequently endorsed by other national bodies and by ISO in Geneva, Switzerland. Its dominance in all fields of scientific and numerical computing grew as new, highly optimizing compilers came onto the market. Although newer languages, particularly Pascal, Basic, PL/1, and later Ada attracted their own adherents, scientific users throughout the 1980s remained true to Fortran. Only towards the end of that decade did C draw increasing support from scientific programmers who had discovered the power of structures and pointers.

During all this time, X3J3 kept functioning, developing the successor version to Fortran 77. It was to be a decade of strife and contention. The early plans, in the late 1970s, were mainly to add to Fortran 77 features that had had to be left out of that standard. Among these were dynamic storage and an array language, enabling it to map directly onto the architecture of supercomputers, then coming onto the market. The intention was to have this new version ready within five years, in 1982. But two new factors became significant at that time. The first was the decision that the next standard should not just codify existing practice, as had largely been the case in 1966 and 1978, but also extend the functionality of the language through

innovative additions (even though, for the array language, there was significant borrowing from John Iverson's APL and from DAP Fortran). The second factor was that X3J3 no longer operated under only American auspices. In the course of the 1980s, the standardization of programming languages came increasingly under the authority of the international body, ISO. Initially this was in an advisory role, but now ISO is the body that, through its technical committee WG5 (in full, ISO/IEC JTC1/SC22/WG5), is responsible for determining the course of the language. WG5 also steers the work of the development body, then as now, the highly skilled and competent X3J3. As we shall see, this shift in authority was crucial at the most difficult moment of Fortran 90's development.

The internationalization of the standards effort was reflected in the welcome given by X3J3 to six or seven European members; they, and about one-third of X3J3's U.S. members, provided the overlapping core of membership of X3J3 and WG5 that was vital in the final years in bringing the work to a successful conclusion. X3J3 membership, which peaked at about 45, is restricted to one voting member per organization, and significant decisions require a majority of two-thirds of those voting. Nationality plays no role, except in determining the U.S. position on an international issue. Members, who are drawn mainly from the vendors, large research laboratories, and academia, must be present or represented at two-thirds of all meetings in order to retain voting rights.

In 1980, X3J3 reported on its plans to the forerunner of WG5 in Amsterdam, Holland. Fortran 8x, as it was dubbed, was to have a basic array language, new looping constructs, a bit data type, data structures, a free source form, a mechanism to "group" procedures, and another to manage the global name space. Old features, including COMMON, EQUIVALENCE, and the arithmetic-IF, were to be consigned to a so-called obsolete module, destined to disappear in a subsequent revision. This was part of the "core plus modules" architecture, for adding new features and retiring old ones, an aid to backwards compatibility. Even though Fortran 77 compilers were barely available, the work seemed well advanced and the mood was optimistic. Publication was intended to take place in 1985. It was not to be.

One problem was the sheer number of new features that were proposed as additions to the language, most of them worthwhile in themselves but with the totality being too large. This became a recurrent theme throughout the development of the standard. One example was the suggestion of Lawrie Schonfelder (Liverpool University), at a WG5 meeting in Vienna, Austria, in 1982, that certain features already proposed as additions could be combined to provide a full-blown derived data type facility, thus providing Fortran with abstract data types. This idea was taken up by X3J3 and has since come to be recognized, along with the array language, as one of the two main advances brought about by what became Fortran 90. However, the ramifications go very deep: all the technical details of how to handle arrays of objects of derived types that in turn have array components that have the pointer attribute, and so forth, have to be precisely defined and rigorously specified.

## **Conflict**

The meetings of X3J3 were often full of drama. Most compiler vendors were represented as a matter of course but, for many, their main objective appeared to be to maintain the status quo and to ensure that Fortran 90 never saw the light of

day. One vendor's extended (and much-copied) version of Fortran 77 had virtually become an industry standard, and it saw as its mission the maintenance of this lead. A new standard would cost it its perceived precious advantage. Other large vendors had similar points of view, although those marketing supercomputers were clearly keen on the array language. Most users, on the other hand, were hardly prepared to invest large amounts of their employers' and their own resources in simply settling for a trivial set of improvements to the existing standard. However, as long as X3J3 worked under a simple-majority voting rule, at least some apparent progress could be made, although the underlying differences often surfaced. These were even sometimes between users — those who wanted Fortran to become a truly modern language and those wanting to maintain indefinite backwards compatibility for their billions of lines of existing code.

At a watershed meeting, in Scranton, Pennsylvania, in 1986, held in an atmosphere that sometimes verged on despair, a fragile compromise was reached as a basis for further work. One breakthrough was to weaken the procedures for removing outdated features from the language, particularly by removing no features whatsoever from the next standard and by striking storage association (i.e., `COMMON` and `EQUIVALENCE`) from the list of features to be designated as obsolescent (as they are now known). A series of votes definitively removed from the language all plans to add: arrays of arrays, exception handling, nesting of internal procedures, the `FORALL` statement (now in Fortran 95), and a means to access skew array sections. There were other features on this list that, although removed, were reinstated at later meetings: user-defined operators, operator overloading, array and structure constructors, and vector-valued subscripts. After many more travails, the committee voted, a year later, by 26 votes to 9, to forward the document for what was to become the first of three periods of public comment.

While the document was going through the formal standards bureaucracy and being placed before the public, X3J3 polished it further. X3J3 also prepared procedures for processing the comments it anticipated receiving from the public, and to each of which, under the rules, it would have to reply individually. It was just as well. Roughly 400 replies flooded in, many of them very detailed and, disappointingly for those of us wanting a new standard quickly, unquestionably negative towards our work. For many it was too radical, but many others pleaded for yet more modern features, such as pointers.

Now the committee was deadlocked. Given that a document had already been published, any further change required not a simple but a two-thirds majority. The conservatives and the radicals could each block a move to modify the draft standard, or to accept a revised one for public review — and just that happened, in Champagne-Urbana, Illinois, in 1988. Any change, be it on the one hand to modify the list of obsolescent features, to add the pointers or bit data type wanted by the public, to add multi-byte characters to support Kanji and other non-European languages or, on the other hand, to emasculate the language by removing modules or operator overloading, and hence abstract data types, to name but some suggestions, none of these could be done individually or collectively in a way that would achieve consensus. I wrote:

"In my opinion, no standard can now emerge without either a huge concession by the users to the vendors (`MODULE / USE`) and/or a major change in the composition of the committee. I do not see how members who have worked for up to a decade



or more, devoting time and intellectual energy far beyond the call of duty, can be expected to make yet more personal sacrifices if no end to the work is in sight, or if that end is nothing but a travesty of what had been designed and intended as a modern scientific programming language. . . . I think the August meeting will be a watershed — if no progress is achieved there will be dramatic resignations, and ISO could even remove the work from ANSI, which is failing conspicuously in its task."

(However, the same notes began with a quotation from *The Taming of the Shrew*: "And do as adversaries do in law, / Strive mightily, but eat and drink / as friend." That we always did, copiously.)

## **Resolution**

The "August meeting" was, unexpectedly, imbued with a spirit of compromise that had been so sadly lacking at the previous one. Nevertheless, after a week of discussing four separate plans to rescue the standard, no agreement was reached. Now the question seriously arose: Was X3J3 incapable of producing a new Fortran standard for the international community, doomed to eternal deadlock, a victim of ANSI procedures?

Breakthrough was achieved at a traumatic meeting of WG5 in Paris, France, a month later. The committee spent several extraordinary days drawing up a detailed list of what *it* wanted to be in Fortran 8x. Finally, it set X3J3 an ultimatum that was unprecedented in the standards world: The ANSI committee was to produce a new draft document, corresponding to WG5's wishes, within five months! Failing that, WG5 would assume responsibility and produce the new standard itself.

This decision was backed by the senior U.S. committee, X3, which effectively directed X3J3 to carry out WG5's wishes. And it did! The following November, it implemented most of the technical changes, adding pointers, bit manipulation intrinsic procedures, and vector-valued subscripts, and removing user-defined elemental functions (now in Fortran 95). The actual list of changes was much longer. X3J3 and WG5, now collaborating closely, often in gruelling six-day meetings, spent the next 18 months and two more periods of (positive) public comment putting the finishing touches to what was now called Fortran 90, and it was finally adopted, after some cliff-hanging votes, for forwarding as a U.S. and international standard on April 11, 1991, in Minneapolis, Minnesota.

Among the remaining issues that were decided along the way were whether pointers should be a data type or be defined in terms of an attribute of a variable, implying strong typing (the latter was chosen), whether the new standard should coexist alongside the old one rather than definitively replace it (it coexisted for a while in the U.S., but was a replacement elsewhere, under ISO rules), and whether, in the new free source form, blanks should be significant (fortunately, they are).

## **Fortran 90**

The main new features of Fortran 90 are, first and foremost, the array language and abstract data types. The first is built on whole array operations and assignments, array sections, intrinsic procedures for arrays, and dynamic storage. It was designed with optimization in mind. The second is built on modules and module procedures, derived data types, operator overloading and generic interfaces, together with

pointers. Also important are the new facilities for numerical computation including a set of numeric inquiry functions, the parametrization of the intrinsic types, new control constructs — `SELECT CASE` and new forms of `DO`, internal and recursive procedures and optional and keyword arguments, improved I/O facilities, and many new intrinsic procedures. Last but not least are the new free source form, an improved style of attribute-oriented specifications, the `IMPLICIT NONE` statement, and a mechanism for identifying redundant features for subsequent removal from the language. The requirement on compilers to be able to identify, for example, syntax extensions, and to report why a program has been rejected, are also significant. The resulting language is not only a far more powerful tool than its successor, but a safer and more reliable one too. Storage association, with its attendant dangers, is not abolished, but rendered unnecessary. Indeed, experience shows that compilers detect errors far more frequently than before, resulting in a faster development cycle. The array syntax and recursion also allow quite compact code to be written, a further aid to safe programming.

No programming language can succeed if it consists simply of a definition (witness Algol 68). Also required are robust compilers from a wide variety of vendors, documentation at various levels, and a body of experience. The first Fortran 90 compiler appeared surprisingly quickly, in 1991, especially in view of the widely touted opinion that it would be very difficult to write one. Even more remarkable was that it was written by one person, Malcolm Cohen of NAG, in Oxford, U.K. There was a gap before other compilers appeared, but now they exist as native implementations for almost all leading computers, from the largest to PCs. For the most part, they produce very efficient object code; where, for certain new features, this is not the case, work is in progress to improve them.

The first book, *Fortran 90 Explained*, was published by John Reid and me shortly before the standard itself was published. Others followed in quick succession, including excellent texts aimed at the college market. At the time of writing there are at least 19 books in English and 22 in various other languages: Chinese, Dutch, French, Japanese, Russian, and Swedish. Thus, the documentation condition is fulfilled.

The body of experience, on the other hand, has yet to be built up to a critical size. Teaching of the language at college level has only just begun. However, I am certain that this present volume will contribute decisively to a significant breakthrough, as it provides models not only of the numerical algorithms for which previous editions are already famed, but also of an excellent Fortran 90 style, something that can develop only with time. Redundant features are abjured. It shows that, if we abandon these features and use new ones in their place, the appearance of code can initially seem unfamiliar, but, in fact, the advantages become rapidly apparent. This new edition of *Numerical Recipes* stands as a landmark in this regard.

## ***Fortran Evolution***

The formal procedures under which languages are standardized require them either to evolve or to die. A standard that has not been revised for some years must either be revised and approved anew, or be withdrawn. This matches the technical pressure on the language developers to accommodate the increasing complexity both of the problems to be tackled in scientific computation and of the underlying hardware

on which programs run. Increasing problem complexity requires more powerful features and syntax; new hardware needs language features that map onto it well.

Thus it was that X3J3 and WG5, having finished Fortran 90, began a new round of improvement. They decided very quickly on new procedures that would avoid the disputes that bedevilled the previous work: WG5 would decide on a plan for future standards, and X3J3 would act as the so-called development body that would actually produce them. This would be done to a strict timetable, such that any feature that could not be completed on time would have to wait for the next round. It was further decided that the next major revision should appear a decade after Fortran 90 but, given the somewhat discomfiting number of requests for interpretation that had arrived, about 200, that a minor revision should be prepared for mid-term, in 1995. This should contain only “corrections, clarifications and interpretations” and a very limited number (some thought none) of minor improvements.

At the same time, scientific programmers were becoming increasingly concerned at the variety of methods that were necessary to gain efficient performance from the ever-more widely used parallel architectures. Each vendor provided a different set of parallel extensions for Fortran, and some academic researchers had developed yet others. On the initiative of Ken Kennedy of Rice University, a High-Performance Fortran Forum was established. A coalition of vendors and users, its aim was to produce an ad hoc set of extensions to Fortran that would become an informal but widely accepted standard for portable code. It set itself the daunting task of achieving that in just one year, and succeeded. Melding existing dialects like Fortran D, CM Fortran, and Vienna Fortran, and adopting the new Fortran 90 as a base, because of its array syntax, High-Performance Fortran (HPF) was published in 1993 and has since become widely implemented. However, although HPF was designed for data parallel codes and mainly implemented in the form of directives that appear to non-HPF processors as comment lines, an adequate functionality could not be achieved without extending the Fortran syntax. This was done in the form of the PURE attribute for functions — an assertion that they contain no side effects — and the FORALL construct — a form of array assignment expressed with the help of indices.

The dangers of having diverging or competing forms of Fortran 90 were immediately apparent, and the standards committees wisely decided to incorporate these two syntactic changes also into Fortran 95. But they didn't stop there. Two further extensions, useful not only for their expressive power but also to access parallel hardware, were added: elemental functions, ones written in terms of scalars but that accept array arguments of any permitted shape or size, and an extension to allow nesting of WHERE constructs, Fortran's form of masked assignment. To readers of *Numerical Recipes*, perhaps the most relevant of the minor improvements that Fortran 95 brings are the ability to distinguish between a negative and a positive real zero, automatic deallocation of allocatable arrays, and a means to initialize the values of components of objects of derived data types and to initialize pointers to null.

The medium-term objective of a relatively minor upgrade has been achieved on schedule. But what does the future hold? Developments in the underlying principles of procedural programming languages have not ceased. Early Fortran introduced the concepts of expression abstraction ( $X=Y+Z$ ) and later control expression (e.g., the DO loop). Fortran 77 continued this with the if-then-else, and Fortran 90 with the DO and SELECT CASE constructs. Fortran 90 has a still higher level of expression abstraction (array assignments and expressions) as well as data structures and even

full-blown abstract data types. However, during the 1980s the concept of objects came to the fore, with methods bound to the objects on which they operate. Here, one particular language, C++, has come to dominate the field. Fortran 90 lacks a means to point to functions, but otherwise has most of the necessary features in place, and the standards committees are now faced with the dilemma of deciding whether to make the planned Fortran 2000 a fully object-oriented language. This could possibly jeopardize its powerful, and efficient, numerical capabilities by too great an increase in language complexity, so should they simply batten down the hatches and not defer to what might be only a passing storm? At the time of writing, this is an open issue. One issue that is not open is Fortran's lack of in-built exception handling. It is virtually certain that such a facility, much requested by the numerical community, and guided by John Reid, will be part of the next major revision. The list of other requirements is long but speculative, but some at the top of the list are conditional compilation, command line argument handling, I/O for objects of derived type, and asynchronous I/O (which is also planned for the next release of HPF). In the meantime, some particularly pressing needs have been identified, for the handling of floating-point exceptions, interoperability with C, and allowing allocatable arrays as structure components, dummy arguments, and function results. These have led WG5 to begin processing these three items using a special form of fast track, so that they might become optional but standard extensions well before Fortran 2000 itself is published in the year 2001.

## **Conclusion**

Writing a book is always something of a gamble. Unlike a novel that stands or falls on its own, a book devoted to a programming language is dependent on the success of others, and so the risk is greater still. However, this new *Numerical Recipes in Fortran 90* volume is no ordinary book, since it comes as the continuation of a highly successful series, and so great is its significance that it can, in fact, influence the outcome in its own favor. I am entirely confident that its publication will be seen as an important event in the story of Fortran 90, and congratulate its authors on having performed a great service to the field of numerical computing.

*Geneva, Switzerland*  
*January 1996*

Michael Metcalf

# License Information

Read this section if you want to use the programs in this book on a computer. You'll need to read the following Disclaimer of Warranty, get the programs onto your computer, and acquire a Numerical Recipes software license. (Without this license, which can be the free "immediate license" under terms described below, the book is intended as a text and reference book, for reading purposes only.)

## ***Disclaimer of Warranty***

**We make no warranties, express or implied, that the programs contained in this volume are free of error, or are consistent with any particular standard of merchantability, or that they will meet your requirements for any particular application. They should not be relied on for solving a problem whose incorrect solution could result in injury to a person or loss of property. If you do use the programs in such a manner, it is at your own risk. The authors and publisher disclaim all liability for direct or consequential damages resulting from your use of the programs.**

## ***How to Get the Code onto Your Computer***

Pick one of the following methods:

- You can type the programs from this book directly into your computer. In this case, the *only* kind of license available to you is the free "immediate license" (see below). You are not authorized to transfer or distribute a machine-readable copy to any other person, nor to have any other person type the programs into a computer on your behalf. We do not want to hear bug reports from you if you choose this option, because experience has shown that *virtually all* reported bugs in such cases are typing errors!
- You can download the Numerical Recipes programs electronically from the Numerical Recipes On-Line Software Store, located at our Web site (<http://www.nr.com>). They are packaged as a password-protected file, and you'll need to purchase a license to unpack them. You can get a single-screen license and password immediately, on-line, from the On-Line Store, with fees ranging from \$50 (PC, Macintosh, educational institutions' UNIX) to \$140 (general UNIX). Downloading the packaged software from the On-Line Store is also the way to start if you want to acquire a more general (multiscreen, site, or corporate) license.
- You can purchase media containing the programs from Cambridge University Press. Diskette versions are available in IBM-compatible format for machines running Windows 3.1, 95, or NT. CDROM versions in ISO-9660 format for PC, Macintosh, and UNIX systems are also available; these include both Fortran and C versions (as well as versions in Pascal

and BASIC from the first edition) on a single CDROM. Diskettes purchased from Cambridge University Press include a single-screen license for PC or Macintosh only. The CDROM is available with a single-screen license for PC or Macintosh (order ISBN 0 521 576083), or (at a slightly higher price) with a single-screen license for UNIX workstations (order ISBN 0 521 576075). Orders for media from Cambridge University Press can be placed at 800 872-7423 (North America only) or by email to [orders@cup.org](mailto:orders@cup.org) (North America) or [trade@cup.cam.ac.uk](mailto:trade@cup.cam.ac.uk) (rest of world). Or, visit the Web sites <http://www.cup.org> (North America) or <http://www.cup.cam.ac.uk> (rest of world).

## ***Types of License Offered***

Here are the types of licenses that we offer. Note that some types are automatically acquired with the purchase of media from Cambridge University Press, or of an unlocking password from the Numerical Recipes On-Line Software Store, while other types of licenses require that you communicate specifically with Numerical Recipes Software (email: [orders@nr.com](mailto:orders@nr.com) or fax: 617 863-1739). Our Web site <http://www.nr.com> has additional information.

- [“Immediate License”] If you are the individual owner of a copy of this book and you type one or more of its routines into your computer, we authorize you to use them on that computer for your own personal and noncommercial purposes. You are not authorized to transfer or distribute machine-readable copies to any other person, or to use the routines on more than one machine, or to distribute executable programs containing our routines. This is the only free license.
- [“Single-Screen License”] This is the most common type of low-cost license, with terms governed by our Single Screen (Shrinkwrap) License document (complete terms available through our Web site). Basically, this license lets you use Numerical Recipes routines on any one screen (PC, workstation, X-terminal, etc.). You may also, under this license, transfer pre-compiled, executable programs incorporating our routines to other, unlicensed, screens or computers, providing that (i) your application is noncommercial (i.e., does not involve the selling of your program for a fee), (ii) the programs were first developed, compiled, and successfully run on a licensed screen, and (iii) our routines are bound into the programs in such a manner that they cannot be accessed as individual routines and cannot practicably be unbound and used in other programs. That is, under this license, your program user must not be able to use our programs as part of a program library or “mix-and-match” workbench. Conditions for other types of commercial or noncommercial distribution may be found on our Web site (<http://www.nr.com>).
- [“Multi-Screen, Server, Site, and Corporate Licenses”] The terms of the Single Screen License can be extended to designated groups of machines, defined by number of screens, number of machines, locations, or ownership. Significant discounts from the corresponding single-screen

prices are available when the estimated number of screens exceeds 40. Contact Numerical Recipes Software (email: [orders@nr.com](mailto:orders@nr.com) or fax: 617 863-1739) for details.

- [“Course Right-to-Copy License”] Instructors at accredited educational institutions who have adopted this book for a course, and who have already purchased a Single Screen License (either acquired with the purchase of media, or from the Numerical Recipes On-Line Software Store), may license the programs for use in that course as follows: Mail your name, title, and address; the course name, number, dates, and estimated enrollment; and advance payment of \$5 per (estimated) student to Numerical Recipes Software, at this address: P.O. Box 243, Cambridge, MA 02238 (USA). You will receive by return mail a license authorizing you to make copies of the programs for use by your students, and/or to transfer the programs to a machine accessible to your students (but only for the duration of the course).

## ***About Copyrights on Computer Programs***

Like artistic or literary compositions, computer programs are protected by copyright. Generally it is an infringement for you to copy into your computer a program from a copyrighted source. (It is also not a friendly thing to do, since it deprives the program’s author of compensation for his or her creative effort.) Under copyright law, all “derivative works” (modified versions, or translations into another computer language) also come under the same copyright as the original work.

Copyright does not protect ideas, but only the expression of those ideas in a particular form. In the case of a computer program, the ideas consist of the program’s methodology and algorithm, including the necessary sequence of steps adopted by the programmer. The expression of those ideas is the program source code (particularly any arbitrary or stylistic choices embodied in it), its derived object code, and any other derivative works.

If you analyze the ideas contained in a program, and then express those ideas in your own completely different implementation, then that new program implementation belongs to you. That is what we have done for those programs in this book that are not entirely of our own devising. When programs in this book are said to be “based” on programs published in copyright sources, we mean that the ideas are the same. The expression of these ideas as source code is our own. We believe that no material in this book infringes on an existing copyright.

## ***Trademarks***

Several registered trademarks appear within the text of this book: Sun is a trademark of Sun Microsystems, Inc. SPARC and SPARCstation are trademarks of SPARC International, Inc. Microsoft, Windows 95, Windows NT, PowerStation, and MS are trademarks of Microsoft Corporation. DEC, VMS, Alpha AXP, and ULTRIX are trademarks of Digital Equipment Corporation. IBM is a trademark of International Business Machines Corporation. Apple and Macintosh are trademarks of Apple Computer, Inc. UNIX is a trademark licensed exclusively through X/Open

Co. Ltd. IMSL is a trademark of Visual Numerics, Inc. NAG refers to proprietary computer software of Numerical Algorithms Group (USA) Inc. PostScript and Adobe Illustrator are trademarks of Adobe Systems Incorporated. Last, and no doubt least, Numerical Recipes (when identifying products) is a trademark of Numerical Recipes Software.

## ***Attributions***

The fact that ideas are legally “free as air” in no way supersedes the ethical requirement that ideas be credited to their known originators. When programs in this book are based on known sources, whether copyrighted or in the public domain, published or “handed-down,” we have attempted to give proper attribution. Unfortunately, the lineage of many programs in common circulation is often unclear. We would be grateful to readers for new or corrected information regarding attributions, which we will attempt to incorporate in subsequent printings.



# Contents

<b><i>Preface to Volume 2</i></b>	<b><i>viii</i></b>
<b><i>Foreword by Michael Metcalf</i></b>	<b><i>x</i></b>
<b><i>License Information</i></b>	<b><i>xvii</i></b>
<b><i>21 Introduction to Fortran 90 Language Features</i></b>	<b><i>935</i></b>
21.0 Introduction	935
21.1 Quick Start: Using the Fortran 90 Numerical Recipes Routines	936
21.2 Fortran 90 Language Concepts	937
21.3 More on Arrays and Array Sections	941
21.4 Fortran 90 Intrinsic Procedures	945
21.5 Advanced Fortran 90 Topics	953
21.6 And Coming Soon: Fortran 95	959
<b><i>22 Introduction to Parallel Programming</i></b>	<b><i>962</i></b>
22.0 Why Think Parallel?	962
22.1 Fortran 90 Data Parallelism: Arrays and Ininsics	964
22.2 Linear Recurrence and Related Calculations	971
22.3 Parallel Synthetic Division and Related Algorithms	977
22.4 Fast Fourier Transforms	981
22.5 Missing Language Features	983
<b><i>23 Numerical Recipes Utility Functions for Fortran 90</i></b>	<b><i>987</i></b>
23.0 Introduction and Summary Listing	987
23.1 Routines That Move Data	990
23.2 Routines Returning a Location	992
23.3 Argument Checking and Error Handling	994
23.4 Routines for Polynomials and Recurrences	996
23.5 Routines for Outer Operations on Vectors	1000
23.6 Routines for Scatter with Combine	1002
23.7 Routines for Skew Operations on Matrices	1004
23.8 Other Routines	1007
<b><i>Fortran 90 Code Chapters</i></b>	<b><i>1009</i></b>
<b><i>B1 Preliminaries</i></b>	<b><i>1010</i></b>
<b><i>B2 Solution of Linear Algebraic Equations</i></b>	<b><i>1014</i></b>
<b><i>B3 Interpolation and Extrapolation</i></b>	<b><i>1043</i></b>