Dear Sir,

Complexity Analyses of Event Set Algorithms
[C. M. Reeves, *The Computer Journal* **27**, 72 (1984)]

I found Professor Reeves' paper concluding that a ternary heap is the data structure to be preferred for the handling of future events in simulations extremely interesting. As I chose to use a binary heap in the Pascal visual simulation system I developed I am also close to his position in practice. However, his analysis is not the end of the story. While considering the best structure for another, simpler, system in BASIC I compared a linked list, a binary heap and exhaustive search of a linear array. Ignoring the linked list in the interests of the most striking comparison I was surprised by the result. In certain circumstances the exhaustive search won by a factor of two!

The reason is that there is another important consideration, ignored in published comparisons, and that is the simultaneity in the event list. By this I mean the average number of events returning at a time beat. Given that one's simulation executive structure is correct (a larger assumption than it should be) then choice of time unit to encourage simultaneous returns diminishes the number of time beats and thus executive overhead time. Such a choice therefore is to be encouraged provided that logical time ordering of events is preserved where necessary. The experiment I conducted was of a simulation with 4096 entries potentially in the event list. I varied the simultaneity between 16 events per time beat and 200 per time beat. The heap was three times as fast in the 16 case whereas the exhaustive search was twice as fast in the 200 case. The reason is obvious as the exhaustive scan finds all the simultaneous events in one pass through the array. I think the size of the experiment and the magnitude of the differences observed is sufficient to establish the point.

Yours faithfully
JOHN G. CROOKES
Department of Operational Research
Lancaster University
UK

Book Reviews

R. N. NEEDHAM AND A. J. HERBERT
The Cambridge Distributed Computing System
Addison-Wesley, London, 1982. 170 pp. £8.50.

'Distributed' computing systems come in many forms, some more distributed than others. That described here incorporates two important basic notions—the Cambridge Ring which provides reliable and simple extensible communications basis, and the interconnection via the ring of special purpose servers dedicated to the provision of specific support functions.

The Cambridge Ring is an empty slot ring in which small packets circulate, providing a nominal 10 Mbits/second data rate between any two points on the ring. The protocol overheads mean that practically sustainable data rates are very much slower. As important as the ring's speed is its high speed switching capability, which permits a single source on the ring to access a multiplicity of destinations with effectively zero call set-up times. The authors describe the evolution of protocols for working across the ring, and then in a loosely historical sequence, they describe the simpler 'servers' which inter-communicate. Amongst the most basic is the file-server, which acts as the single long-term repository for all users' data. Other servers provide for program loading and initiation (in still other servers), for terminal concentration, printing, error logging and time stamping. An important simple server, the 'name server', provides a mapping between a named server and its physical address, thus avoiding some of the problems which arise if a service fails—or is for any reason transferred from one server to another.

Probably the most important material is provided by the three chapters which describe the concept of a processor bank, together with its associated management and scheduling. Rather than providing processing capability by giving access to a single large system running a multi-user environment, the approach adopted is to provide a 'bank' of smaller systems, each run as a single-user system. These small systems are not all the same, and each is offered as effectively a naked processor plus memory, all other support being accessed via the ring. In the atmosphere of a research laboratory, with dominant interests in operating systems, and with attendant interests in resource scheduling, control, and the management of access to both processors and stored data this is clearly a challenging approach, providing the maximum freedom to individuals.

This is a short book. There is much that it does not cover, most noticeably a discussion of the extent to which the ideas it presents can form the basis of general systems in which the individual users are much less concerned with the construction of systems, and perhaps much less tolerant of some of its idiosyncracies. In my view, it is an important book, and should be bought by all those with an interest in the development of computing systems. Having bought the book, they should read it.

M. WELLS
Leeds

P. C. RAVASIO, G. HOPKINS AND N. NAFFAH (EDITORS)
Local Computer Networks
North-Holland, Amsterdam. 555 pp. $65.00.

This book contains reprints of the papers presented at an IFIP Symposium on Local Computer Networks held in April 1982. A total of 33 papers is presented, under 10 different headings, although in one or two cases it is not entirely clear what relevance the paper has to the section in which it appears.

The editors have done a commendable job in actually collecting the texts of the papers from the speakers. As is inevitable, the level of presentation is uneven, ranging from detailed theoretical analyses of the stability of phase-locked loops, or the behaviour of queueing systems, to rather general appeals of the motherhood/apple-pie variety. Obviously there is little that editors can do about this. Less forgiveable is the absence of any report of any discussions during the Symposium. There is an index of contributors names, but no subject index.

The content is variable, and suffers from that curious form of time-warp, which leaves the reader uncertain as to whether an article describes a system which is in existence and working, or is little more than a design study. The main emphasis is on the hardware for the transport of bits, and on its overall performance, and some of the material here is of a high standard, well presented and authoritative. There is some attention to higher-level protocols and of linking networks; I would have welcomed more.

Overall, this is a book for the library, rather than one's own bookshelf.

M. WELLS
Leeds

J. MARTIN
Strategic Data-Planning Methodologies
Prentice-Hall, Englewood Cliffs, New Jersey, 1982. 236 pp. £26.00.

The publishers of this book inform us, the readers, that the author has written 28 major books on computers and telecommunications. Your reviewer is conscious that he could only think of 24 titles and of these he could only claim to be familiar with about six. One criticism often levelled at James Martin is that his books are repetitive. Your reviewer would support this argument but point out that it is probably a deliberate act by the author who

does not wish to make the mistake of trying to divide the world of computing into water-tight compartments. Martin's books cover overlapping areas and are designed to make the reader aware of the close relationship that exists between these neighbouring technologies.

This latest book covers the Data Base world and concentrates on the need for careful planning either when introducing data base techniques or simply contemplating changes to existing systems. It is impossible to quarrel very seriously with anything Martin writes. He is always careful to avoid controversy and instead concentrates on such universally acceptable proposals such as 'Implementing the whole of an information architecture will take years. It is necessary to decide what to do first . . .'. One is tempted to add that a similar discipline is required by the potential reader of Martin's library of Computing books.

It is very difficult to know what market this book will serve. Like all successful authors, Martin has a following and no doubt libraries will wish to keep their collections up to date. Busy data processing managers may not find it very inviting and the paucity of supporting evidence for Martin's assertions will not endear the book to academic circles. It will be interesting to see what 1983 offers the James Martin fan.

J. O. JENKINS
London

M. JAMBU AND M-O. LEBEAUX
Cluster Analysis and Data Analysis
North-Holland, Amsterdam, 1983. 888 pp. Dfl. 200.00.

This is a formidable tome which is divided into two parts. Part 1, 'Methods and algorithms', occupies 254 pages and Part 2, 'Computer programs and user's guide', fills the remainder. At the end of each part there is a list of references, over 300 in the case of Part 1 and just 20 for Part 2.

The authors use Cluster Analysis to mean 'the set of computer procedures which build up or recognise hierarchies or partitions underlying data sets'. By Data Analysis they mean those procedures which 'describe, recognise or identify structures underlying clouds of points' e.g. multidimensional scaling, factor analysis and related techniques.

The treatment is mathematical throughout and some original material is included in Part 1. The programs presented in Part 2 are written in FORTRAN; there are 40 programs and 326 subroutines representing a total of 25 000 statements.

The text includes some examples of how the programs are used but they would not be of much help to the newcomer to the subject. In many ways this is a reference book and it is not intended that it should be read from start to finish. Rather the researcher would select the sections of interest. Unfortunately there is no index and thus the table of contents must be relied upon. Although this is quite detailed it does not really serve as an adequate index and so it is difficult to find the sections of interest.

This book is, as one would expect, very up-to-date and since it is so comprehensive it can be recommended as a reference volume. However, as a guide to *users* of Cluster Analysis this reviewer believes that it is quite incomprehensible and in this context it does not replace Sokal and Sneath's *Principles of numerical taxonomy* which remains the most readable authoritative work. The translation from the original French is excellent and full credit for this must go to Ms H. Teil.

G. N. LANCE
Bristol

NORMA KLISS LEHMKUHL
Fortran 77 A Top-Down Approach
Macmillan, New York, 1983. 466 pp.

This is an introductory text intended for students and assumes no background other than elementary algebra: it explicitly does not aim to be a reference book on Fortran. The order in which language concepts are introduced indicates the orientation of the book: first list-directed output of character constants, then assignment for real and integer variables, further list-directed input and output, control structures and formatted input–output. These are followed by variable types other than real and integer and by one- and two-dimensional arrays. The book is almost finished before subprograms and common blocks appear. A final chapter summarizes auxiliary input–output statements including OPEN and CLOSE, and an appendix covers some of the remaining statements such as ENTRY, EQUIVALENCE, and SAVE. Some commonly used statements like BLOCK DATA, IMPLICIT, INQUIRE and unformatted input–output are not described; nor are most of the new possibilities for character manipulation available in Fortran 77 compared with Fortran 66.

There is strong emphasis throughout the book on first defining an algorithm to solve the given problem, on drawing a flowchart to represent the algorithm and on coding from the flowchart. There are numerous exercises at the end of each chapter. The language described is Fortran 77; some, but not all, deviations from the ISO standard are noted, the most noticeable unremarked divergence being the description of the OPEN statement.

The text takes the student through the language very slowly, very thoroughly and very competently up to 30 pages from the end of the last chapter. It comes as a surprise that there is a confused description of common blocks and an incorrect interpretation of adjustable arrays. Nevertheless the book would make a suitable text for a student wanting to learn the language for general interest or for use as an occasional tool. Other introductory texts would be more appropriate for anyone proposing to earn their living wholly or partly by writing Fortran.

DAVID MUXWORTHY
Edinburgh

J. D. FOLEY AND A. VAN DAM
Fundaments of Interactive Computer Graphics
Addison-Wesley, London, 1982. 664 pp. £15.95.

The two authors of this book have been major contributors to the area of computer graphics over the last decade. This book captures a great deal of their knowledge and is a major addition to the computer graphics literature. The emphasis is on interactive systems. The book is divided into four main sections with the first section giving a complete introduction to the subject (both hardware and software) using an interactive office layout system as the application to be defined and a simple graphics package in which to describe it. Particular attention is paid to the concepts needed for good interactive design. The algorithms required to code the graphics package (coordinate transformations, clipping, segment storage, etc.) are given in detail. All the examples are coded in PASCAL and given in full.

The second section concentrates on the mathematics of viewing in three dimensions including perspective. It introduces the models required to hold the application data and how they relate to the graphical data structures. The third section describes display files and other hardware assists given by modern display technology. Particular attention is paid to raster displays with algorithms for line drawing, anti-aliasing, area filling and polygon clipping.

The final section is the most original, concentrating on the production of realistic synthetic images. Hidden surface removal, shading, texturing, lighting and colouring are all described. Most of the new work in this area is either described or referenced.

The authors are to be congratulated on a major achievement. Beautifully produced with exciting illustrations, it is recommended to anybody with an interest in computer graphics.

F. R. A. HOPGOOD
Didcot