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