for $\theta = 0{\cdot}8$ illustrates the high utilisation approximation of Section 3.3. At lower utilisations the curves wilt progressively. The broken curve shows the result of applying equations (3.37, 3.39) outside their range of validity with $x = 0{\cdot}7$—i.e. above the critical value 0·6263 of equation (3.23). The existence of a turning point and subsequent negative gradient is behaviour clearly incompatible with that of a respectable cumulative distribution function.

4. Discussion

On the personal level, quite the most satisfying part of the work reported here was the Eureka experience of stumbling, after some weeks of ineffectual groping, upon the realisation that generating functions were the key to the problem. Not only did this technique permit a compact and elegant formulation of the equations, but, in the case of unit length requests, their solution then required merely schoolboy algebra. Although the introduction of this technique into the study of storage allocation may possibly be the most important contribution of this paper, the prediction of a stability threshold at 37% store utilisation is undoubtedly the most intriguing. Any speculation now, in advance of further investigation, is likely to be a source of embarrassment in the future but some comment is clearly called for.

The theoretical model has two distinct components. The major part is summed up in equations (2.14, 2.39) which lead to equilibrium conditions in terms of the parameters p_2 and k. The second is somewhat separate and is contained in Section 2.3.2. This establishes a relationship between p_2 and k so that the free store profiles $\tilde{\phi}$ form a one-parameter family. It is a simple matter to transform between k, the block utilisation ratio, as parameter and θ, the store utilisation. As was noted at the time, the argument of Section 2.3.2, though logically sound, is somewhat unsatisfying. The difficulty is a familiar one in that a slightly different but equally plausible interpretation of randomness might yield a different relationship. Hopefully the present information-theoretic argument will be wholly replaced by a specific analysis based upon the problem situation. This will in all probability lead to a different equation (2.34) and hence either to a revised value for the stability threshold or, if equation (3.19) is then satisfied throughout $0 \leqslant \theta \leqslant 1$, to a disappearance of the phenomenon entirely.

The availability of a closed analytical solution for unit sized requests simplifies further investigation of the sensitivity of $\tilde{\phi}$ to the relationship between p_2 and k. Results from simulations are required for comparison and it is hoped to report conclusions in due course. Of particular interest will be empirical results from the critical region of 37% store utilisation.

Also left for future investigation is the solution of the equilibrium equations for general request distributions. This represents a non-trivial extension on which there appears to be little published data from simulations apart from hints that free store patterns are relatively insensitive to the request distribution.

In the shadow of all this unfinished business, some may consider the present publication premature. Others will agree that these preliminary findings are of interest and may stimulate parallel activity.

5. Acknowledgements

Most of the ideas that have been described were thought through during a recent period of convalescence. I gratefully acknowledge both the expertise of the doctors whose good offices culminated in securing me this pseudo-sabbatical, and the willingness of my departmental colleagues to undertake my normal duties in the same period. In the preparation of the final manuscript I have been greatly helped by the insight and advice of the referee, and I offer him my sincere thanks.

References

CAMPBELL, J. A. (1971). A note on an optimal-fit method for dynamic allocation of storage, *The Computer Journal*, Vol. 14, pp. 7-9.

CRANSTON, B. and THOMAS, R. (1975). A simplified recombination scheme for the Fibonacci buddy system, *CACM*, Vol. 18, pp. 331-332.

GELENBE, E., BOEKHORST, J. C. A. and KESSELS, J. L. W. (1973). Minimising wasted space in partitioned segmentation, *CACM*, Vol. 16, pp. 343-349.

HINDS, J. A. (1975). An algorithm for locating adjacent storage blocks in the buddy system, *CACM*, Vol. 18, pp. 221-222.

HIRSCHBERG, D. S. (1973). A class of dynamic memory allocation algorithms, *CACM*, Vol. 16, pp. 615-618.

ISODA, S., GOTO, E. and KIMURA, I. (1971). An efficient bit-table technique for dynamic storage allocation of 2^n-word blocks, *CACM*, Vol. 14, pp. 589-592.

KNOWLTON, K. C. (1965). A fast storage allocator, *CACM*, Vol. 8, pp. 623-625.

KNUTH, D. E. (1968). *The art of computer programming*, Vol. 1: Fundamental algorithms, (1st edn. Addison-Wesley) §2.5 Dynamic storage allocation, pp. 435-455.

NIELSEN, N. R. (1977). Dynamic memory allocation in computer simulation, *CACM*, Vol. 20, pp. 864-873.

PURDOM, P. W. and STIGLER, S. M. (1970). Statistical properties of the buddy system, *JACM*, Vol. 17, pp. 683-697.

RANDELL, B. (1969). A note on storage fragmentation and program segmentation, *CACM*, Vol. 12, pp. 365-369, 372.

SHEN, K. K. and PETERSON, J. L. (1974, 1975). A weighted buddy method for dynamic storage allocation, *CACM*, Vol. 17, pp. 558-562 and Vol. 18, p. 202.

SHORE, J. E. (1977). Anomalous behaviour of the fifty-percent rule in dynamic memory allocation, *CACM*, Vol. 20, pp. 812-820.

WOLMAN, E. (1965). A fixed optimum cell-size for records of various lengths, *JACM*, Vol. 12, pp. 53-70.

Book review

The Process of Question Answering, by Wendy G. Lehnert, 1978; 278 pages. (*Wiley*, £11·75)

This book seeks to answer the sixty-four dollar question: What is a question and what is an answer? Although a laudable attempt is made to present a systematic and logical flowchart for this common and complex human activity, the book suffers considerably from that tiresome American scholarly affliction, a combination of indigestible jargon and fourth-form humour. An example of the first runs as follows: 'The processes which assign the proper conceptual category to a question must be sensitive to the context in which that question is asked. Questions cannot be correctly understood by processes that do not consider contextual factors.' In other words, a question like 'Are you a little queer?' depends on its context if it is to generate the correct response. An example of the humour is one subsection heading which goes: 'Smart heuristics know when to quit'. Another subsection heading, which falls between these stools, unintentionally bursts into verse: 'You can't always expect to find / Exactly what you had in mind'. No wonder one of the programs exploited in the production of this book is called QUALM.

R. W. LAST (Hull)