A Lumped-state Model of Clustering in Dynamic Storage Allocation

C. M. Reeves
Department of Computer Science, The University, Keele, Staffs ST5 5BG, UK

A fragmentation model based upon the concept of an achievable equilibrium has proved successful with large systems but breaks down when the numbers of reserved blocks or free fragments are small. A Markov chain model is presented which, by assuming thorough spatial mixing and a strict temporal alternation of reservations and releases, reduces computational complexity to manageable proportions in systems of intermediate size. Good continuity across the two models is demonstrated.

1. INTRODUCTION

A recent paper[1] has described results obtained from a particular model of free-store fragmentation in which expected changes in configuration resulting from the block allocation and release mechanisms were derived. In equilibrium, allocations and releases take place at the same mean rate and the equilibrium configuration was determined as that for which the combined effects of allocations and releases produce no net change.

Formally the model yields two loosely coupled sets of equations. One set determines the manner in which reserved blocks cluster together in contiguous sequences of varying size. The other determines the distribution of sizes of the fragments of free store which separate the clusters. The two sets are linked by two simple equations. One has been well known for many years as the fifty per cent rule.[2] This has the form

$$\frac{F}{B} = \tfrac{1}{2}p \tag{1}$$

where F and B are the numbers of free fragments and of reserved blocks in the store, and p is the probability that a request for space will be allocated from an oversize free fragment. It is convenient to write

$$\frac{2F}{B} = x \tag{2}$$

and then the fifty per cent rule may be expressed as

$$p = x \tag{3}$$

The second equation is new and determines the proportion of isolated reservations (i.e. clusters having just a single member). Denoting by c_r the number of clusters of r reservations, the relation is

$$\frac{c_1}{B} = \frac{x^2}{(x+4)} \tag{4}$$

or equivalently, from the definition of x,

$$\frac{c_1}{F} = \frac{2x}{(x+4)} \tag{5}$$

The fragmentation equations involve the size distribution of reservations but the cluster equations, being concerned with block counts rather than block sizes, do not. In this sense the cluster analysis has the greater generality.

The predictions of the model are in general consistent and satisfactory. At high levels of utilization the composition of the store is insensitive to the absolute sizes of B and F, depending upon their ratio as determined by x. At lower B and F values however, the model breaks down. Indeed, in this region the cluster equations yield negative values for some of the c_r counts!

The most likely explanation lies in the assumption that the mean store configuration is achievable. This is a reasonable and justified approximation when configurations are dense around the mean but at low occupation numbers it is equivalent to the indiscretion, familiar to statisticians, of looking for the average family with 2.4 children.

This then is the background to the present paper. We recognize that the equilibrium should be treated as dynamic. The appropriate model is a Markov chain and this is set up in Section 2 for the analysis of clustering. This paper concentrates upon clustering, partly because of its generality and partly because it was here that the imperfections of the earlier model were first revealed. A comparable treatment of fragmentation remains a possibility for the future.

It quickly becomes apparent that the number of configurations which are, in principle, candidates for recognition as separate states increases dramatically with B. Various devices for controlling this explosion are described in Sections 3 and 4. Two of them deserve mention here.

Rather than treat each spatial arrangement of a given set of cluster sizes as a separate state, the assumption of thorough spatial mixing allows them to be treated as one single state. The significance of this lumping together of configurations is reflected in the title of this paper.

Secondly, we recall the substantial success of the earlier model at higher B values. There, the achievability of the mean configuration suggests that transitions are concentrated within a cloud of states close to the mean. The more distant states have low probability of occurrence and are thus not significant. The program organization described in Section 4 exploits this concept to identify dynamically, and to record, only those states which are judged to be significant.

CCC-0010-4620/84/0027-0135$04.00

2. THE MARKOV MODEL

A store configuration consists of clusters of contiguous reservations separated by blocks of free store, the whole regarded as circular so that end effects may be disregarded. The size of a cluster is a count of the number of reservations it contains: we are not concerned with the amount of storage occupied.

The specific approximation introduced by the present model is that a system state is characterized by a set of counts c_r, $r \geq 1$, where c_r is the number of clusters of size r. All orderings of clusters around the store are equivalent and the model assumes random mixing. Thus the probability that a randomly selected cluster is of size r is c_r/F where F is the number of clusters, and hence of free blocks, given by

$$F = \sum_r c_r \tag{6}$$

Similarly the number B of reservations is given by

$$B = \sum_r r c_r \tag{7}$$

Transitions between states take place through the two basic mechanisms of reservation and release. In the one case, all or part of a sufficiently large free block is allocated to accommodate a fresh reservation and in the other, the space occupied by an existing reservation is freed.

2.1 The reservation mechanism

The model assumes the so-called random allocation strategy in which a sufficiently large free block is selected at random. In a real situation one can envisage the possibility that heavy fragmentation may combine with a large request size so that there is no way of satisfying the request. Similarly it is possible either that the largest fragment exactly matches the request or that all free fragments are larger than the current request size. These are extreme possibilities corresponding to the certainty of an exact fit and of an oversize fit, respectively. More usually one may expect a range of fragment sizes, some less than, some equal to and some greater than the size of the current request. Each of the fragments in the latter two categories is then equally likely to be picked to accommodate the request.

The present model assumes that a sufficiently large free block can always be found and that the probability p of selecting an oversize fragment is constant. This is the only point at which the physical sizes of store blocks enter the discussion. In modelling terms this represents a powerful simplification. Its justification rests upon the expectation that the dominant configurations will have broadly similar fragmentation patterns.

When an oversize free block is used, the new reservation is allocated space at one end rather than in the middle in order to avoid needless fragmentation of free store. The probability that the adjacent cluster which it joins has size r initially is pc_r/F and the effect on the size counts is $c_r \leftarrow c_{r-1}$ and $c_{r+1} \leftarrow c_{r+1} + 1$.

The probability of an exact fit is $q = 1 - p$. If the sizes of the neighbouring clusters on the left and right are r and s respectively, then the effect is to yield a single cluster of size $r + s + 1$. The effect on the size counts is $c_r \leftarrow c_r - 1$, $c_s \leftarrow c_s - 1$ and $c_{r+s+1} \leftarrow c_{r+s+1} + 1$. The associated probability may be written as $q\Psi_{rs}$ where

$$\Psi_{rs} = \frac{c_r(c_r - 1)}{F(F-1)}, \quad \text{if } r = s$$

$$= \frac{c_r}{F}\frac{c_s}{(F-1)}, \quad \text{if } r \neq s \tag{8}$$

This expression depends upon the assumption of thorough spatial mixing of clusters. The left-hand cluster, say, is randomly selected and then the right-hand one is selected independently at random from those that remain.

2.2 The release mechanism

The model makes the usual assumption of random releases. Each current reservation is equally likely to be released next. Three categories arise, according as the released block has 0, 1 or 2 free neighbours.

Category 0 blocks are the internal members of a cluster of size r, $r \geq 3$. Let the reservations be numbered 0 to $r - 1$. The effect of releasing the block in position u, $1 \leq u \leq r - 2$, is to split the cluster into two, having sizes u and $r - u - 1$. Thus, with probability c_r/B the effect on the size counts for each of the $(r - 2)$ possible values of u is $c_r \leftarrow c_r - 1$, $c_u \leftarrow c_u + 1$ and $c_{r-u-1} \leftarrow c_{r-u-1} + 1$.

Category 1 blocks are the end members of a cluster of size r, $r \geq 2$. The space freed is merged with the adjacent free block and the cluster size reduces to $r - 1$. Thus with probability $2c_r/B$ the effect is $c_r \leftarrow c_r - 1$ and $c_{r-1} \leftarrow c_{r-1} + 1$.

The category 2 blocks are the lone reservations: the clusters of size 1. The space released is merged with the two adjacent free blocks. Thus with probability c_1/B the effect is $c_1 \leftarrow c_1 - 1$.

2.3 Equilibrium

These mechanisms enable the probabilities to be determined of a transition from any cluster configuration of B blocks to each configuration of $B + 1$ blocks in the case of a reservation and to each configuration of $B - 1$ blocks in the case of a release. In statistical equilibrium, reservations and releases take place at the same mean rate and so preserve a mean level of store utilization.

The author's earlier paper treated mean cluster counts as actual values. Equations were formed expressing the requirement that the net effect of reservations and releases on the equilibrium set of counts is zero. This assumption is satisfactory when the counts are sufficiently large but was found to be suspect at low store loadings.

Betteridge[3] has given an analysis in terms of a Markov system in which each distinct spatial configuration of blocks is regarded as a separate state. This approach is rigorous but leads to equations of immense complexity and size. It seems likely that its application is limited to small systems having $B \leq 10$ say.

The present work seems to occupy the middle ground between these extremes. The assumption of thorough spatial mixing enables the various permutations of a particular partitioning of B reservations into F clusters to be lumped together into a single effective state. This

too is an approximation and we shall explore its consequences.

Ideally one would like to include all B-values in the reckoning, with transitions which are equally likely to be reservations or releases. This would add greatly to the complexity of the calculations, and there are unresolved theoretical difficulties about selecting a dominant B value to determine mean store loading.

Instead therefore we regard each transition as a two-stage process, a release followed by a reservation, applied only to states with a specified B value. There is thus an intermediate state corresponding to $B-1$ reservations. This mechanism is chosen for its simplicity. We could equally well have chosen a reservation followed by a release, but this would be more complicated since there are many more configurations of $B+1$ intermediate states than of $B-1$. Similarly, more complex combinations were also excluded. With the model as proposed, therefore, there are two free parameters, B and p.

Suppose that the different c-sets corresponding to the different states of a system of B reservations are ordered in some way and let X_r be the equilibrium probability of the rth state. Similarly, order the states of a system of $B-1$ reservations and let Y_r be the equilibrium probability of the rth such intermediate state. The mechanisms of reservation and release described above enable us to compute the transition matrices $\mathbf{G}$ and $\mathbf{H}$ where G_{rs} is the probability of a transition from the sth B-state to the rth $(B-1)$-state, and H_{sr} is the probability of the reverse transition. The equilibrium requirement is thus given by the two equations

$$\mathbf{Y} = \mathbf{GX} \tag{9}$$

and

$$\mathbf{X} = \mathbf{HY} \tag{10}$$

The overall transition matrix is

$$\mathbf{T} = \mathbf{HG} \tag{11}$$

and the requirement

$$\mathbf{X} = \mathbf{TX} \tag{12}$$

determines $\mathbf{X}$ as the quantity of theoretical interest. Once $\mathbf{X}$ is available, it is a simple matter to compute the mean values of F and of each c_r.

3. ENUMERATION OF STATES

The states of a system of B reservations are the partitions of B, taking no account of the ordering of components. Each cluster is a component so that the number of components is F. We wish to determine for any given system both how many states there are, and how to order them. The number of states determines the size of the problem and the ordering facilitates referencing states.

It would be convenient to group states according to their F value but this is found to create substantially increased complexity and so has not been pursued.

Denote by $P(B, n)$ the number of partitions of B into components of size n or less. The cases of practical concern correspond to $1 \le n \le B$. It is clear that if $n > B$ then $P(B, n) = P(B, B)$.

In general, the number of partitions of B having r as the largest component is $P(B-r, r)$ since there remain $B-r$ items to be grouped into components still of size r or less. It follows that

$$P(B, n) = P(B, n-1) + P(B-n, n) \tag{13}$$

Applying the obvious boundary conditions, we may calculate $P(B, n)$ as follows

$$\begin{aligned} P(B, n) &= 1, \quad \text{if } n = 1 \\ &= P(B, n-1) + P(B-n, k), \quad \text{if } 1 < n < B \\ &= P(B, n-1) + 1, \quad \text{if } n = B \end{aligned} \tag{14}$$

where

$$k = \min(n, B-n) \tag{15}$$

The total number of partitions for any B is $P(B, B)$. This increases rapidly with B as is shown by the following specimen values:

B	7	10	13	20	30	40	50
$P(B, B)$	15	42	101	627	5604	37 338	204 226

It is clear that the feasibility of solving Eqn (12) depends upon exploiting the special features of the problem to give a more compact representation than that implied by the conventional matrix notation.

3.1 Ordering of states

We turn now to the problem of arranging the partitions for given B into a standard order so that we may reference them by position number w, $1 \le w \le P(B, B)$. It is convenient first to transform from the set $(c_1, c_2, \ldots)$ of counts to the corresponding partition $(p_1, p_2, \ldots)$ with components in decreasing order. Thus for $B = 6$, the state with $c_1 = 3, c_3 = 1$ becomes the partition 3111.

The states are then arranged in the lexicographical order of their partitions. Thus if $(p_1, p_2, \ldots)$ and $(q_1, q_2, \ldots)$ are two such partitions P and Q, then P precedes or follows Q according as $p_i < q_i$ or $p_i > q_i$ where p_i and q_i are the leftmost non-matching pair. For example, the ordering of the 7 partitions of 5 is (11111), (2111), (221), (311), (32), (41), (5).

The position number w of a partition $(p_1, p_2, \ldots)$ is obtained by counting the number of partitions that precede it. First there are $P(B, p_1 - 1)$ whose largest component is less than p_1. Then there are $P(B-p_1, p_2 - 1)$ whose largest component matches p_1 and whose second component is less than p_2. And so on. Thus

$$\begin{aligned} w = 1 &+ P(B, p_1 - 1) + P(B - p_1, p_2 - 1) \\ &+ \cdots + P(B - p_1 - p_2 - \cdots - p_{r-1}, p_r - 1) \end{aligned} \tag{16}$$

where p_{r+1} would be 1 or 0.

Conversely, given w, we may determine the partition and hence the set of counts. Thus p_1 is the smallest x such that $P(B, x) \ge w$. p_2 is the smallest x such that $P(B-p_1, x) \ge w - P(B, p_1 - 1)$, and so on until the partition is complete.

For given B we require algorithms which will convert in either direction between the position number and the count-set representation of states. Recursive evaluation of $P(B, n)$ is very slow and we have preferred to evaluate and pre-store the elements so that they may be obtained by table look-up in a linearized triangular array. For B values up to 50, an array of $50 \times \frac{51}{2} = 1275$ elements is sufficient.

Though time-consuming in execution, these conversions permit a substantial economy in the space required to represent states.

4. THE STATE VECTOR REPRESENTATION

As with every choice of data structure, all depends upon the needs of the operations to be performed upon it. In this case it was decided to solve Eqn (12) by the method of successive multiplication. An initial estimate of the vector **X** of state probabilities is repeatedly multiplied by the transition matrix **T** until convergence is achieved. Each iteration may be considered as consisting of four steps therefore:

$$\mathbf{Z} := \mathbf{X}; \quad \mathbf{Y} := \mathbf{GX}; \quad \mathbf{X} := \mathbf{HY}; \quad \text{compare}(\mathbf{Z}, \mathbf{X}) \qquad (17)$$

Considering the step $\mathbf{Y} := \mathbf{GX}$, the elements X_s may be taken in any convenient order and the contributions $G_{rs}X_s$ to each Y_r accumulated. Similarly with the step $\mathbf{X} := \mathbf{HY}$. Both **G** and **H** are sparse and so the two steps via **Y** are preferred to the explicit use of **T**. However, even though **G** and **H** are sparse, there are too many non-zero elements to store them explicitly and so they are generated afresh as required.

These various economies are not in themselves sufficient. The size of the **X**, **Y** and **Z** vectors is an added problem. Previous experience suggests, particularly at larger B values, that the significant states at equilibrium will have count-sets which lie close to their mean values. It is thus expected that the **X**, **Y** and **Z** vectors will themselves be sparse as regards significant elements.

These considerations lead to the representation of the state vectors as disc files of key-value pairs (r, V_r) ordered on r. Thus, to compute $\mathbf{Y} := \mathbf{GX}$ for example, the range of subscripts of **Y** is split into slices such that a slice of **Y** elements fits into main store. The **X** file is scanned once per slice and the selected **Y** elements are accumulated. Those pairs (r, Y_r) for which $Y_r > \varepsilon$, for some small ε, are written to the file **Y**. The sum of the significant **Y** elements is appended to the file so that the vector **Y** may be renormalized on subsequent use.

The number of states judged insignificant by this criterion is hoped to be large. The parameter ε must be chosen small enough for the net probability of all the insignificant states taken together to be small.

This device is not logically watertight. It is possible that the effect of truncating the state vector could be to confine successive iterates to a region of their domain which excludes the true solution. It is hoped that the consistency and quality of the results, particularly in relation to the limit $B \to \infty$, will persuade the reader, as it does the author, that such fears have not been realized. A pointer in this direction must be the fairly strong connectivity of the transition matrix.

5. IMPLEMENTATION DETAILS

A program was written in Pascal and run in batch mode on the GEC 4082 system at Keele. The main runs were with $\varepsilon = 0.000005$ at intervals of 0.1 in p for a range of B values, namely 7, 10, 13 and 20. It was intended to extend the range of B up to about 50 but it became clear that this would have been more antisocial than the results would have warranted.

A major objective was to study the behaviour of system properties as B increases towards infinity. The usefulness of the very limited range of B values thus depends critically upon their ability to establish a clear trend. In the circumstances it was felt that the niceties of numerical extrapolation would be no more convincing than simple plots of properties against $1/B$. The chosen B values are well spread in the scale of $1/B$ and, of course, having infinity on the paper does facilitate extrapolation!

The precision of real arithmetic on the GEC 4082 corresponds to a little over six significant decimal digits. As all matrix elements and vector components are non-negative, there is no loss of accuracy through differencing in the calculations.

6. THE SIGNIFICANT CONFIGURATIONS

For $B = 7$ and $B = 10$ the calculations with $\varepsilon = 0.000005$ showed that all configurations were significant. At the higher B values, a number of configurations are not significant. Figures 1(a) and 1(b) show the proportions of significant states for $B = 13$ and $B = 20$. These are seen to be sensitive to p, being largest near $p = 0.8$.

The net probability associated with the insignificant states is at most 4ε for $B = 13$ and 28ε for $B = 20$. Not surprisingly, this probability is greatest where the number

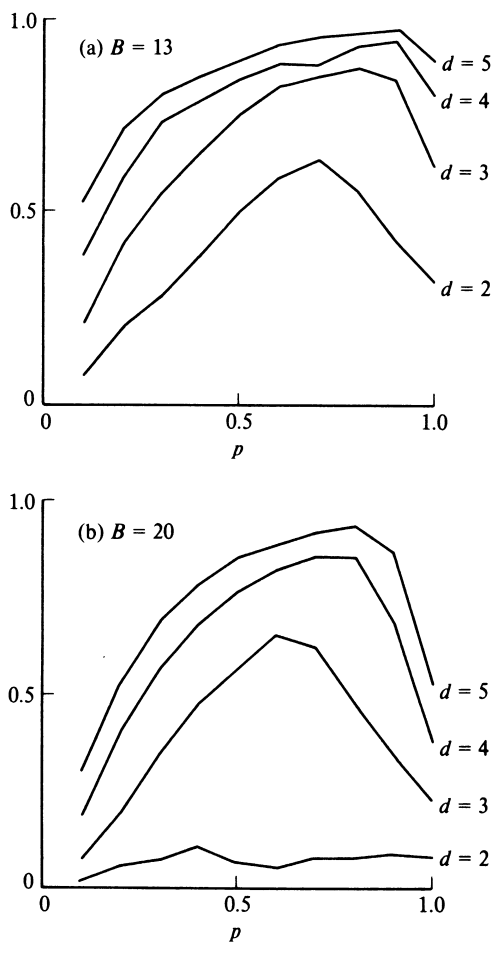

Figure 1. Proportions of significant configurations.

of insignificant states is greatest, i.e. near $p = 0$ and $p = 1$.

Also shown in Figs 1(a) and 1(b) are the effects of varying ε. Writing $\varepsilon = 0.5 \times 10^{-d}$, curves are plotted for $d = 5$, 4, 3 and 2. Those for $d = 5$ are obtained from the calculations as described. Those for $d = 4$, 3 and 2 are obtained by counting the numbers of elements of the solution vector for $d = 5$ which exceed the corresponding threshold.

It is seen that, for any given ε, as B increases the proportion of significant configurations does indeed drop but not at a rate which competes with the increase in the total number of configurations. Thus the complexity of each calculation rises sharply with B, this being most apparent near $p = 0.8$. In practical terms this made it impolitic to extend the main series of calculations beyond $B = 20$ though two isolated runs were performed for $B = 30$ at $p = 0.2$ and 1.0 giving the proportions of significant configurations as 0.29 and 0.16 respectively for $\varepsilon = 0.000005$ with 230ε and 106ε as the probabilities of insignificant states. Knuth pointed out that with a mixture of request sizes, the effective value of p will lie close to 1.0.

Finally in this section it may be of interest to list quite arbitrarily in Table 1 the dozen most probable configurations for the case $B = 20$, $p = 0.6$. Their individual c-count sets may be compared with the mean values derived by averaging over all 558 configurations significant at the $\varepsilon = 0.000005$ level.

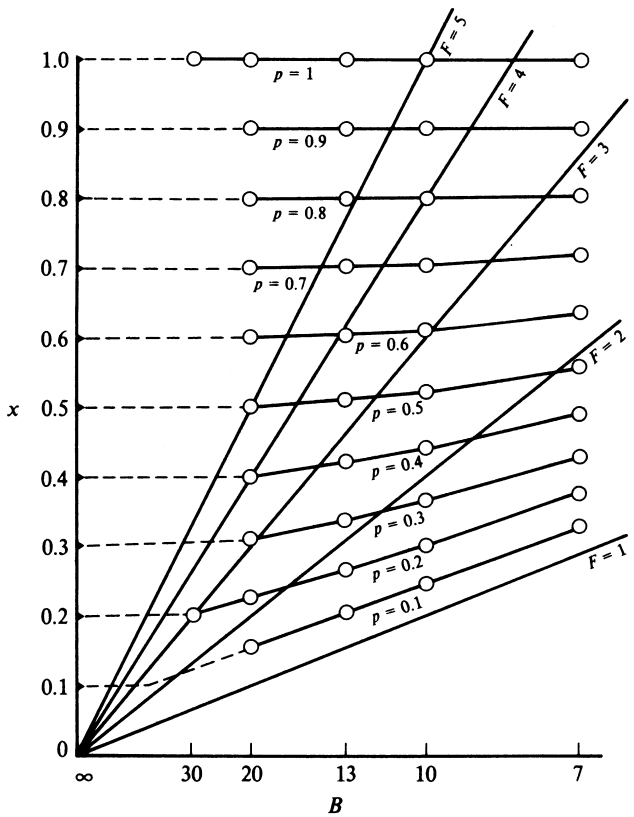

Figure 2. Variation of x with B and p.

Table 1. 12 most significant configurations for $B = 20, p = 0.6$

r	X_r	c_7	c_6	c_5	c_4	c_3	c_2	c_1
149	0.012526			1	1	2	2	1
145	0.010206			1	1	1	3	2
228	0.009138		1		1	2	1	2
157	0.009094			1	2	1	1	2
311	0.008846	1			1	1	2	2
247	0.008832		1	1		1	2	2
226	0.008110		1		1	1	3	1
314	0.008063	1			1	2	1	1
250	0.008056		1	1		2	1	1
236	0.008030		1		2	1	1	1
182	0.008021			2	1	1	1	1
89	0.007973				2	2	2	2
	means	0.211	0.312	0.459	0.676	0.997	1.416	1.485

Table 2. Computed mean x values

p \ B	7	10	13	20	30
0.1	0.327	0.246	0.203	0.154	
0.2	0.374	0.301	0.264	0.225	0.200
0.3	0.428	0.365	0.336	0.309	
0.4	0.489	0.439	0.418	0.402	
0.5	0.558	0.521	0.508	0.500	
0.6	0.634	0.609	0.603	0.599	
0.7	0.717	0.703	0.701	0.699	
0.8	0.806	0.800	0.800	0.799	
0.9	0.901	0.900	0.900	0.900	
1.0	1.000	1.000	1.000	1.000	1.002

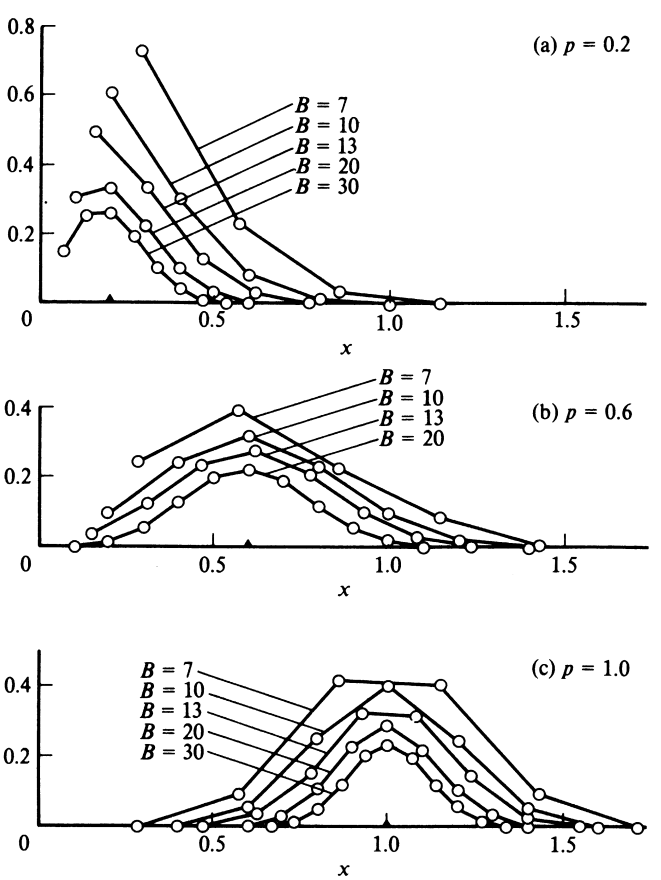

Figure 3. Probabilities of achievable x values.

7. NUMBERS OF FRAGMENTS

By accumulating the probabilities of individual configurations with respect to the number of fragments in each, the probability of each value of F is obtained, and hence the mean value. From this the mean value of $x = 2F/B$ is determined. Table 2 lists these results which are then plotted against $1/B$ for each p in Fig. 2.

In Figs 3(a), 3(b) and 3(c) the probabilities of each x value are plotted for a selection of p and B values. The concentration of x values as B increases is apparent at each p. At low p values, since F is a whole number, values of B less than $2/p$ result in all achievable x values, and hence their mean, being above the putative limiting value p. For each p, as soon as B is big enough to yield a mean F value above about 3 (i.e. $B \geq 6/p$) the x distribution adopts a quasi-normal form with mean close to p.

In Fig. 2 the radial lines correspond to the achievable values of F. Above $F = 3$ good agreement with the fifty per cent rule $x = p$ is observed. Below $F = 3$, x is seen to vary approximately linearly with $1/B$.

8. NUMBERS OF ISOLATED RESERVATIONS

A similar analysis applied to c_1 in place of F yields for each B, p pair a probability for each achievable value of c_1 and a corresponding mean value. Table 3 lists the mean values of c_1/B together with the putative limiting values $p^2/(p+4)$ derived from Eqns (3) and (4). These results are plotted against $1/B$ for each p in Fig. 4.

In Figs 5(a), 5(b) and 5(c) the probabilities of each c_1/B value are plotted for a selection of p and B values.

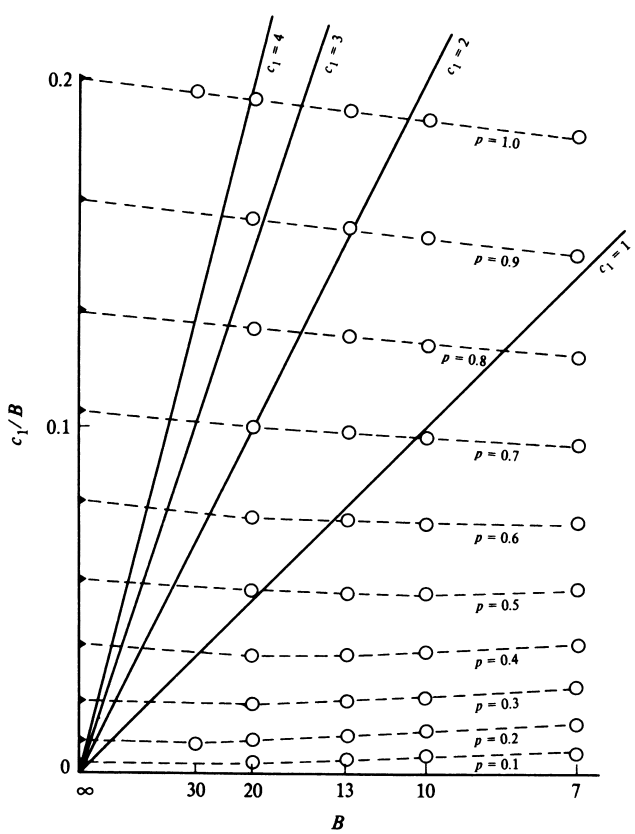

Figure 4. Variation of c_1/B with B and p.

Table 3. Computed mean c_1/B values

p \ B	7	10	13	20	30	$p^2/(p+4)$
0.1	0.006	0.005	0.004	0.003		0.0024
0.2	0.014	0.012	0.011	0.010	0.008	0.0095
0.3	0.025	0.022	0.021	0.020		0.0209
0.4	0.037	0.035	0.034	0.034		0.0364
0.5	0.053	0.052	0.052	0.053		0.0556
0.6	0.072	0.072	0.073	0.074		0.0783
0.7	0.094	0.096	0.098	0.100		0.1043
0.8	0.120	0.123	0.126	0.128		0.1333
0.9	0.149	0.154	0.157	0.160		0.1653
1.0	0.183	0.188	0.191	0.194	0.197	0.2000

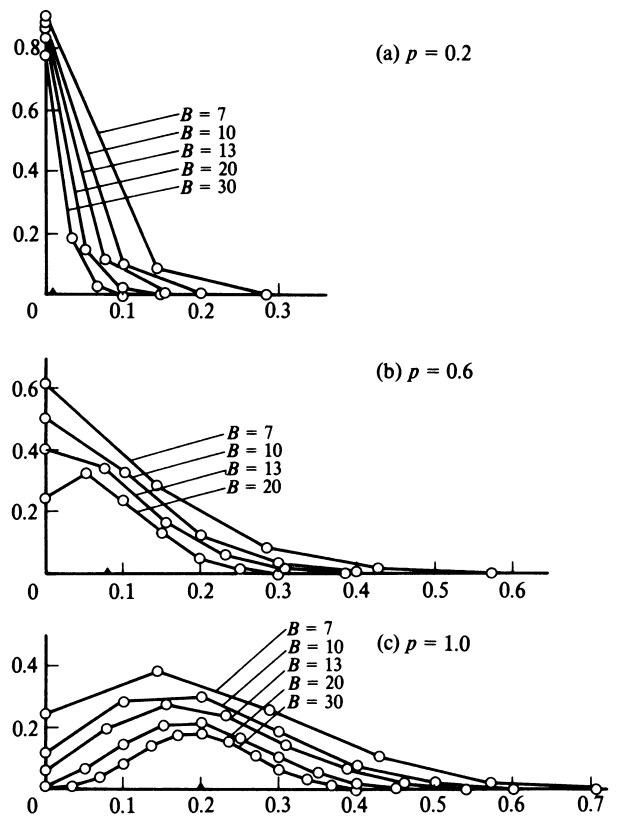

Figure 5. Probabilities of achievable c_1/B values.

As with Fig. 3 an increasing concentration near the limit is observed with increasing B for each p. In this case the achievable values always straddle the limit, since $c_1 = 0$ is achievable. Once again it is seen that for fixed p, and for B large enough to yield three or more achievable values below the limit, the distribution of c_1/B adopts a quasi-normal form.

In Fig. 4 the radial lines correspond to the achievable values of c_1. There is some evidence in Fig. 4 to suggest that, for each p, the mean value of c_1/B approaches the limit from below as B increases. At higher p values this monotonic behaviour persists right through the B range. At lower values of p and of B however, the mean value of c_1/B decreases with increasing B to a turning point below the limiting value. This is clearly seen in Table 3 where for $p = 0.5$ the minimum value is about 0.052 at $B = 11$ and the limiting value is 0.0556. With lower p values the location of the minimum moves towards higher B values.

9. CLUSTER-SIZE PROFILE

The proportion of clusters of size r is $\sigma_r = c_r/F$. Mean values of σ_r for each r should be derived by averaging c_r/F with respect to the probability of each configuration. More conveniently, a pseudo-mean has been calculated as the ratio of the computed mean values of c_r and F. This is wrong in principle but in practice is unlikely to be misleading. In each case the values for each r sum to unity as they should and it is unlikely that significant errors, if present, would so neatly compensate each other.

Table 4 shows values of σ_r for $p = 1.0$. The right-hand column lists the values predicted by the analytical model of the earlier paper. The case $p = 1$ is special in that its predicted σ distribution is independent of B. It is expected therefore that this should correspond to the limit of the present Markov model predictions as B increases to infinity. In Fig. 6 σ_r is plotted against $1/B$ for various r and this expectation is fulfilled in the most gratifying manner. For each r, σ_r is seen to vary linearly with $1/B$ over the whole range of B studied.

Tables 5(a) and 5(b) provide similar details for the case $p = 0.6$. Table 5(a) applies to the Markov model and Table 5(b) to the earlier model where now σ_r depends upon B for $r > 2$. Indeed it was the collapse of that model at low F values that prompted the present study. Figure 7 shows σ_r plotted against $1/B$ for various r with continuous curves corresponding to the Markov model

Table 4. Computed mean σ_r values at $p = 1.0$

r \ B	7	10	13	20	30	∞
	F 3.500	5.000	6.500	10.004	15.032	
1	0.365	0.377	0.382	0.389	0.393	0.4000
2	0.371	0.358	0.352	0.345	0.342	0.3333
3	0.184	0.180	0.178	0.176	0.174	0.1714
4	0.061	0.063	0.064	0.065	0.065	0.0667
5	0.015	0.017	0.018	0.019	0.020	0.0212
6	0.003	0.004	0.004	0.005	0.005	0.0057
7	0.001	0.001	0.001	0.001	0.001	0.0013
8		0.000	0.000	0.000	0.000	0.0001
9		0.000	0.000	0.000	0.000	0.0001

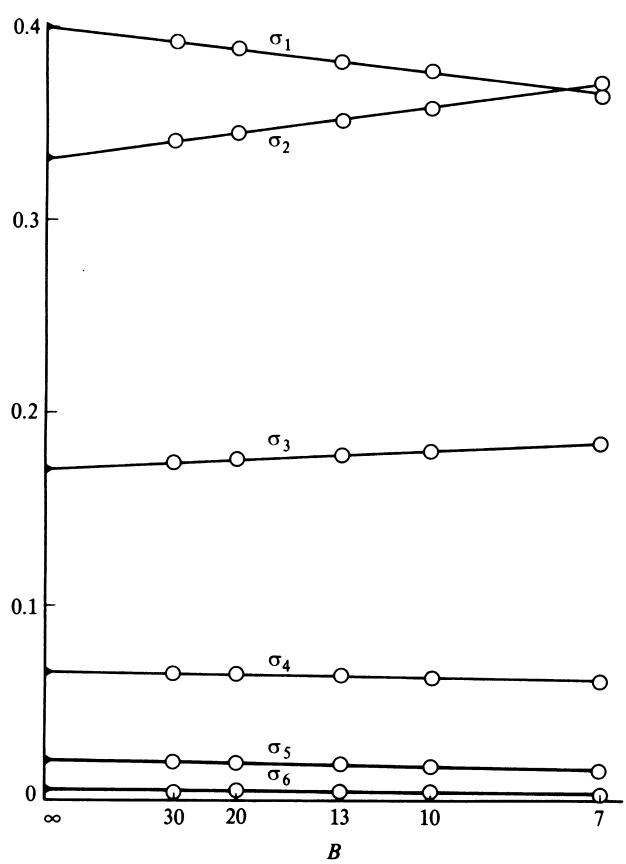

Figure 6. Variation of σ_r with B at $p = 1.0$.

Table 5(a). Computed mean σ_r values at $p = 0.6$

r \ B	7	10	13	20
	F 2.218	3.047	3.917	5.991
1	0.226	0.236	0.242	0.248
2	0.248	0.243	0.240	0.236
3	0.173	0.170	0.169	0.166
4	0.117	0.115	0.114	0.113
5	0.080	0.078	0.077	0.077
6	0.045	0.053	0.052	0.052
7	0.111	0.035	0.035	0.035
8		0.024	0.024	0.024
9		0.014	0.016	0.016
≥ 10		0.032	0.031	0.033

Table 5(b). Analytical σ_r values at = 0.6

r \ B	7	10	20	40	∞
1	0.2609	0.2609	0.2609	0.2609	0.2609
2	0.2308	0.2308	0.2308	0.2308	0.2308
3	0.1381	0.1490	0.1590	0.1599	0.1623
4	0.1240	0.1179	0.1122	0.1117	0.1104
5	0.0668	0.0709	0.0741	0.0744	0.0751
6	0.0647	0.0587	0.0531	0.0525	0.0512
7	0.0347	0.0343	0.0346	0.0347	0.0349
8	0.0328	0.0287	0.0250	0.0246	0.0238
9	0.0139	0.0151	0.0160	0.0160	0.0162
> 10	0.0333	0.0337	0.0343	0.0345	0.0344

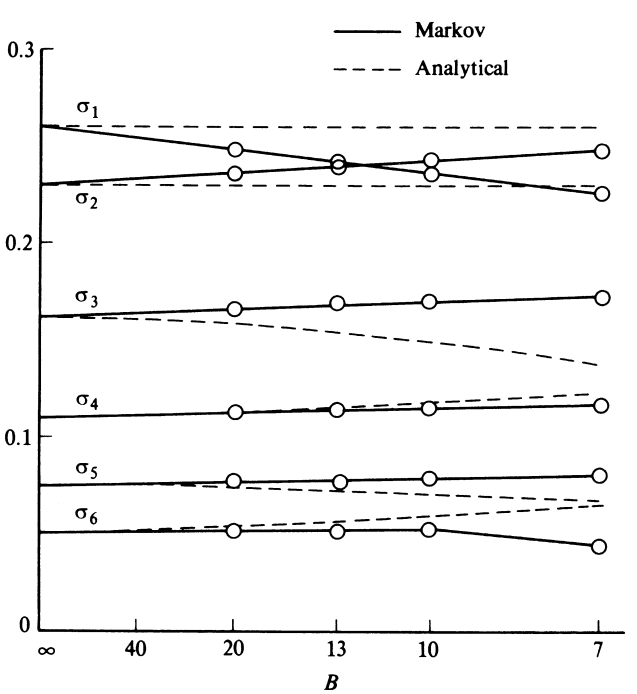

Figure 7. Variation of σ_r with B at $p = 0.6$.

and broken curves to the analytical model. There is again clear indication of agreement between the models as B increases. The Markov model predicts strong linearity of σ_r against $1/B$ for the lower r values whereas there is evident non-linearity in the analytical model.

An interesting speculation raised by these curves is whether the inclusion of B-dependence in the analytical model is worth while. At small B values it causes the model to collapse. At large values it is irrelevant. At intermediate values it yields poorer agreement with the Markov model than does the limiting case. On the other hand, there are no compelling reasons to support the superiority of the Markov model as a reflection of 'the truth'. The strength of the case for including B-dependence in the analytical model is its essential role in the development of the high and low utilization regimes in the analysis of the size distribution of free fragments.

As a final observation, it turns out happily that 20 is quite a good approximation to infinity. We have investigated a number of properties of the system and have concluded that, where there are at least two or three achievable values on either side of the mean, there is good agreement between the two models. Those ancient shepherds with one–two–many as their counting system had a point. In retrospect it seems likely that an earlier appreciation of their wisdom would have led to a different program design. Accepting a more limited range of B, a better balance between space and time resources would come from filing the matrix elements of the transitions rather than regenerating them at each iteration.

REFERENCES

1. C. M. Reeves, Free store distribution under random fit allocation: Part 3. *The Computer Journal* **26**, 25–35 (1983).
2. D. E. Knuth, *The Art of Computer Programming, Vol 1: Fundamental Algorithms.* Addison Wesley (1968).
3. T. Betteridge, An algebraic analysis of storage fragmentation. *Ph.D. thesis*, University of Newcastle upon Tyne (1979).

Received October 1982