Free Store Distribution under Random Fit Allocation: Part 3

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

The analyses of earlier papers are extended to cover general distributions of reservation sizes. Equations are derived for the equilibrium distributions of clusters of reservations and of free fragments, and numerical methods of solution are described. The model is applied in three simple experiments to investigate the influence of key parameters.

1. SUMMARY

This paper corrects and extends a model of the statistical mechanics of dynamic storage allocation that was reported in two previous papers.[1,2] The opportunity is taken to give here a fairly complete review of the model, indicating any significant amendments. Though slightly repetitious, this is felt to be preferable to the alternative of a catalogue of nit-picking cross references.

The review is contained in section 2. It is convenient to discuss first the aggregation of reserved blocks into contiguous groups. The size of a cluster is the number of reservations in it and should not be confused with the amount of store that it occupies. The outcome from this analysis is the set of equations controlling the equilibrium distribution of cluster sizes for various store loadings. Being concerned with block counts rather than block sizes, the equations do not depend upon the size distribution of reservations. The parameters of the equations are B, the number of reservations, and F, the number of free fragments.

Two important corollaries are: firstly the fifty per cent rule[3]

$$p = x \tag{1}$$

where $x = 2F/B$ and p is the probability that in allocating space for a reservation, an oversize free block will be used; secondly the relation

$$\sigma_1 = 2x/(x + 4) \tag{2}$$

where σ_1 is the proportion of clusters of unit size. This latter relation transforms directly into

$$p_2 = x^2/(x + 4) \tag{3}$$

where p_2 denotes the proportion of reserved blocks having two free neighbours, i.e. the proportion of lone reservations.

It should be noted that the cluster analysis in the present section 2 is more general than that in Ref. 2 in that it covers both finite and infinite store sizes.

Equation (3) is important in the analysis of free store fragmentation which then follows. Here we are concerned with the distribution of fragments by size and the sizes of reservations are significant. As with clustering, the analysis applies to both finite and infinite stores. Regrettably there was an error in Ref. 2 in the line labelled (7.6): the denominator should be $(F - 1)$ in place of F. As a consequence the fifty per cent rule is again recovered, thus removing the even greater embarrassment of an apparent conflict with the cluster analysis.

A notable feature of the model is the identification of two distinct regimes characterized by the level of store utilization. At high utilizations fragmentation is complete and the store comprises a macroscopically homogeneous mixture of clusters and free fragments. At low utilizations the free store comprises one long free block and a range of fragments. The two regimes merge together at a critical threshold level of utilization.

In section 3, methods are derived for solving the various model equations. Attention is concentrated upon numerical methods of general application. The cluster equations are straightforward and are solved by direct substitution in the sequence of recurrence relations. The fragmentation equations, on the other hand, require numerical iteration. As presently formulated, the equations are poorly conditioned and the accuracy achieved in the solutions is severely limited, particularly in the interesting threshold region. It may be that better luck or sounder judgment will yield an improved method of solution.

Solutions of the cluster equations are discussed in section 4. It is observed that results are insensitive to storesize and depend primarily upon the utilization as characterized by x. At low B or F values, however, the model breaks down by predicting negative counts for some cluster sizes. It is suggested that this arises because the present model requires high cluster counts at each size in order that mean values may reasonably be treated as actual values. At the lower count levels a more rigorous Markov model[4] is appropriate.

Turning to the fragmentation model, comparisons are made in section 5 with alternative solution methods in the special case of unit-length requests. Here a direct numerical solution is available and, in the special case of a large store at high utilization, an analytic solution. These comparisons indicate the general level of corruption induced by differencing errors in the iterative calculations. Calculations with a nominal precision of eleven significant decimal digits can yield results of no more than three or four reliable digits in a bad case.

With the model duly constructed and validated, the paper continues in section 6 with an account of its use as a predictive tool. Three simple experiments are described, each designed to give fresh insight into the effects of varying the distribution of request sizes.

The first experiment considers requests whose sizes are uniformly distributed in the range 1 to R, for various

CCC-0010-4620/83/0026-0025 $05.50

R. Of particular interest is the location of the boundary between the high and low utilization regimes. Knuth's[3] expectation that x approaches unity as the variability of request sizes increases is verified. The actual utilization at the threshold, measured by the fraction of the store that is occupied, is however insensitive to R and remains at about 45–50%.

The second experiment was intended to test the influence of the variance of the request size distribution. Two sizes were taken, with equal weight and a fixed mean of 8. As the sizes are moved apart, so the variance is increased. It is evident that the case of sizes 6 and 10, say, is the same as for 3 and 5 with a changed unit of length. The interesting result is that the x threshold mainly reflects the factoring of sizes and is insensitive to changes in variance within each factor class. This effect was unexpected and the third experiment was aimed at shedding further light on it.

In the final experiment, fixed sizes 3 and 5 were taken but with variable weights β and $1-\beta$ where $0 \leq \beta \leq 1$. The extreme values correspond to unit-length requests with suitable scaling of the unit of size. It was of interest to see whether the solutions vary continuously with β at the extremes. The evidence obtained suggests that this is the case.

To conclude, it is clear that the ability to compute, for the first time, profiles of the fragmentation and clustering induced by any distribution of request sizes makes a significant advance in the study of dynamic storage allocation. There is room for improvement in the numerical processes which implement the model. There is wide scope for exploring what are the significant questions which the model should be used to answer. The applications described here are simple exercises but they have shown up the importance of exploiting common factors in the set of request sizes. The consistent, albeit fuzzy, location of the θ threshold gives a second version of a fifty per cent rule. Formally, the result embodied in Eqns (2) and (3) is important. That $dx/d\theta = 0$ at the high utilization threshold is an interesting conjecture awaiting proof (see Section 6.1).

Perhaps the major achievement of the model is the recognition of the two regimes corresponding to high and low store utilization. It is interesting to note that in neither regime is the size of the store a significant parameter. Consider an experiment, starting with an empty large finite store, in which the rate of reservations slightly exceeds the rate of releases. The store slowly fills, maintaining the equilibrium profiles at each stage. We take up the story when some fifty or more reservations are in the store so that our model is applicable.

Because the store is taken to be circular, the reserved blocks and smaller free fragments form a more or less homogeneous mixture which occupies a contiguous region. The rest of the store is one long free block. As the store fills, the long free block is eroded to provide both space for the added reservations and additional free fragments. The mix of reservations and fragments scarcely changes in composition.

Eventually the long free block shrinks so that it can no longer be distinguished from the other free blocks. One can make an analogy with liquid–vapour equilibrium in an enclosed vessel which is heated. The long block is the liquid and the rest is vapour. We have reached the point at which all the liquid is vaporized: the dew point. The dew point is the threshold where we pass from the low to the high utilization regime.

As the store continues to fill, space for more reservations can only be found by reducing the number and size of free fragments. The character of the distributions thus changes, though the mixture retains its macroscopic homogeneity. At some point before the store is full it becomes impossible to satisfy a request for space from the depleted stock of fragments. Our model then fails since we have assumed an unsaturated store. Beyond this point, in the region having low fragment counts, the Betteridge approach would seem more suitable.

2. THE ANALYTICAL MODEL

This section reviews the earlier analysis and extends it where necessary to cover general request-size distributions. Conditions are obtained for statistical equilibrium between reservations and releases in a store of N words. In order to avoid consideration of the effects at the two ends of a linear store, the store is assumed to be circular. The store is also assumed to be unsaturated so that each request for space may be satisfied immediately.

The distribution of request sizes is assumed known: b_r is the probability that a request is for a block of r words. The size of the largest request is denoted by R. In equilibrium, reservations and releases take place at the same mean rate and maintain a store utilization θ, i.e. θ is the fraction of the store which is reserved. Under the conditions studied, requests, current reservations and releases each have the same size distribution.

The model treats the random-fit allocation strategy. Each block of free store of sufficient size is equally likely to be used to satisfy a request. The probability that an oversize free block is selected is denoted by p. In this case the reservation is made at one end of the free block, leaving the residue free for subsequent use. Each end of an oversize block is equally likely to be used.

Similarly with releases: it is assumed that each current reservation is equally likely to be released next. The space occupied is immediately freed and any contiguous free blocks are compacted to form a single larger fragment.

The analysis considers both the clustering of reservations into contiguous sequences and the fragmentation of free store into blocks of various sizes. The average number of reservations is denoted by B and of free blocks by F. Under the assumption of a circular store, the number of clusters of reservations is also F. An important parameter of the system is x where

$$x = 2F/B \tag{4}$$

2.1 The clustering of reservations

Throughout this section it is the counts of blocks that are significant rather than their sizes. In particular the results obtained do not depend upon the request-size distribution.

Clustering is described by the generating function

$$\tilde{c}(a) = \sum_r c_r a^r \tag{5}$$

where c_r, $r \geq 1$, is the number of clusters of r contiguous reservations. It follows that

$$\tilde{c}(1) = \sum_r c_r = F \tag{6}$$

and

$$\tilde{c}'(1) = \sum_r r c_r = B \tag{7}$$

The probability that a random cluster is of size r is denoted by σ_r where

$$\sigma_r = c_r/F \tag{8}$$

The corresponding generating function $\tilde{\sigma}$ defined by

$$\tilde{\sigma}(a) = \sum_r \sigma_r a^r \tag{9}$$

satisfies

$$\tilde{\sigma}(1) = 1 \tag{10}$$

and

$$\tilde{\sigma}'(1) = 2/x \tag{11}$$

2.1.1 Reservations. Two cases arise. In the case where an oversize fragment is used to satisfy a request, the cluster adjacent to the new reservation is extended by one block. Clusters are randomly distributed and so, with probability $p\sigma_r$ for each r, $c_r \leftarrow c_r - 1$ and $c_{r+1} \leftarrow c_{r+1} + 1$.

In the event of an exact fit, the probability that the left and right adjacent clusters are of sizes r and s is ψ_{rs} where

$$\begin{aligned}\psi_{rs} &= \frac{c_r}{F}\frac{c_s}{F-1} = \frac{F}{F-1}\sigma_r\sigma_s \quad \text{if } r \neq s \\ &= \frac{c_r}{F}\frac{c_r - 1}{F-1} = \frac{F}{F-1}\sigma_r\left(\sigma_r - \frac{1}{F}\right) \quad \text{if } r = s\end{aligned} \tag{12}$$

It follows that, with probability $(1-p)\psi_{rs}$ for each r and s,

$$c_r \leftarrow c_r - 1, \quad c_s \leftarrow c_s - 1 \quad \text{and} \quad c_{r+s+1} \leftarrow c_{r+s+1} + 1$$

The expected increment $\Delta^+\tilde{c}$ to $\tilde{c}$ following a single reservation is thus given by

$$\begin{aligned}\Delta^+\tilde{c} = {} & p\sum_r \sigma_r(a^{r+1} - a^r) \\ & + (1-p)\sum_r\sum_s \psi_{rs}(a^{r+s+1} - a^r - a^s)\end{aligned} \tag{13}$$

which simplifies to give

$$\begin{aligned}\Delta^+\tilde{c} = {} & a(1-p)\tilde{\sigma}^2 + (ap + p - 2)\tilde{\sigma} \\ & + \frac{a(1-p)}{F-1}(\tilde{\sigma}^2 - \tilde{\sigma}(a^2))\end{aligned} \tag{14}$$

2.1.2 Releases. Three cases arise corresponding to the degree of compaction required.

(i) *No compaction.* Suppose the released block is in position u in a cluster of length r, $r > 2$, where $1 \leq u \leq r-2$ and where the two end blocks are regarded as in positions 0 and $r-1$. The effect of the release is to split the cluster into two, having lengths u and $r-u-1$. Thus, since each reservation is equally likely to be released, the probability is c_r/B for each r and u that $c_r \leftarrow c_r - 1$, $c_u \leftarrow c_u + 1$ and $c_{r-u-1} \leftarrow c_{r-u-1} + 1$.

(ii) *Single compaction.* Suppose the released block is at one end of a cluster of length r, $r \geq 2$. Then with probability $2c_r/B$ for each r, $c_r \leftarrow c_r - 1$ and $c_{r-1} \leftarrow c_{r-1} + 1$.

(iii) *Double compaction.* Suppose the released block forms a cluster of length 1, i.e. it is a lone block. Then with probability c_1/B the effect is $c_1 \leftarrow c_1 - 1$.

Combining these results, the expected increment $\Delta^-\tilde{c}$ *to* $\tilde{c}$ following a single release is

$$\begin{aligned}\Delta^-\tilde{c} = \frac{1}{B}\Bigg\{ & \sum_{r>2} c_r \sum_{u=1}^{r-2}(a^u + a^{r-u-1} - a^r) \\ & + \sum_{r\geq 2} 2c_r(a^{r-1} - a^r) - c_1 a\Bigg\}\end{aligned} \tag{15}$$

Now

$$\sum_{u=1}^{r-2}(a^u + a^{r-u-1} - a^r) = \frac{2}{1-a}(a - a^{r-1}) - (r-2)a^r \tag{16}$$

so that, after some manipulation,

$$\Delta^-\tilde{c} = \frac{x}{2(1-a)}\{2(a - \tilde{\sigma}) - a(1-a)\tilde{\sigma}'\} \tag{17}$$

2.1.3 Equilibrium. With reservations and releases taking place at the same mean rate, the equilibrium condition is

$$\Delta^+\tilde{c} + \Delta^-\tilde{c} = 0 \tag{18}$$

2.1.4 The fifty per cent rule. Setting $a = 1$ in the above condition yields

$$\Delta^+F + \Delta^-F = 0 \tag{19}$$

thus expressing the constancy of the number of free blocks. On substituting the expressions found for $\Delta^+\tilde{c}$ and $\Delta^-\tilde{c}$, it follows readily (see Part 2, section 2.3.1) that

$$p = x \tag{20}$$

which is equivalent to the fifty per cent rule.

2.1.5 The cluster equations. Combining the general equilibrium condition with the fifty per cent rule to eliminate p yields the cluster equation

$$\begin{aligned}ax(1-a)\tilde{\sigma}' - {} & 2a(1-a)(1-x)\tilde{\sigma}^2 \\ & + 2(2 - 2a + a^2x)\tilde{\sigma} - 2ax \\ & - \frac{2a(1-a)(1-x)}{F-1}(\tilde{\sigma}^2 - \tilde{\sigma}(a^2)) = 0\end{aligned} \tag{21}$$

This equation is to be regarded as an identity in a. It is convenient to set out the separate equations corresponding to the different powers of a. Defining S_n, the coefficient of a^n in $\tilde{\sigma}^2(a)$ by

$$S_n = \sum_r \sigma_r\sigma_{n-r} \tag{22}$$

and picking out the coefficient of a^n in the cluster equation, it follows that for $n \geq 1$,

$$\begin{aligned}(nx + 4)\sigma_n - {} & ((n-1)x + 4)\sigma_{n-1} + 2x\sigma_{n-2} \\ & - 2x\delta_{n1} - 2(1-x)\left(\frac{F}{F-1}\right)(S_{n-1} - S_{n-2}) \\ & + \frac{2(1-x)}{F-1}(\sigma_{(n-1)/2} - \sigma_{(n-2)/2}) \\ = {} & 0\end{aligned} \tag{23}$$

where δ is the Kronecker delta ($\delta_{ij} = 1$ if $i = j$, 0 if $i \neq j$) and $\sigma_i = 0$ if i is non-integral.

2.1.6 Lone reservations. The special case $n = 1$ of the above equations is of particular significance, leading to the relation

$$\sigma_1 = \frac{2x}{x+4} \tag{24}$$

It should be noted that this, being independent of F, extends the analysis of Part 2, which, by taking ψ_{rs} to be $\sigma_r\sigma_s$, effectively assumed F infinite.

2.2 The fragmentation of free store

Both simulation studies and the earlier analyses have shown a significant difference between the free-store fragmentation patterns when operating in conditions of high or low utilization. At high utilizations the proportion of fragments of given length decreases steadily with length, whereas at low utilizations the fragments include one long free block.

Fragmentation is described by the generating function

$$\tilde{f}(a) = \sum_r f_r a^r \tag{25}$$

where f_r, $r \geq 1$, is the number of free fragments of size r words. By convention the long fragment is omitted from the sum at low utilization so that

$$\begin{aligned} \tilde{f}(1) = \sum_r f_r = F, &\quad \text{at high utilization} \\ = F - 1, &\quad \text{at low utilization} \end{aligned} \tag{26}$$

The probability that a random fragment is of size r is denoted by ϕ_r where

$$\phi_r = f_r/F \tag{27}$$

The corresponding generating function $\tilde{\phi}$ defined by

$$\tilde{\phi}(a) = \sum_r \phi_r a^r \tag{28}$$

satisfies

$$\begin{aligned} \tilde{\phi}(1) = 1, &\quad \text{high} \\ = 1 - 1/F, &\quad \text{low} \end{aligned} \tag{29}$$

and the probability at low utilization that a random fragment is the long one is $1/F$.

The mean fragment size (excluding the long one) is denoted by $\bar{r}$. At high utilization it is given by

$$\bar{r} = \sum_r r\phi_r = \tilde{\phi}'(1) \tag{30}$$

At low utilizations, because the long fragment is excluded,

$$\bar{r} = \sum_r r f_r/(F-1) = \frac{F}{F-1}\tilde{\phi}'(1) \tag{31}$$

2.2.1 Reservations. The probability that a request for size i will be met from a free fragment of size r, $r \geq i$, is denoted by Q_{ir}. Thus

$$Q_{ir} = \phi_r/\gamma_i \tag{32}$$

where

$$\begin{aligned} \gamma_i = \sum_{r \geq i} \phi_r, &\quad \text{high} \\ = \sum_{r \geq i} \phi_r + 1/F, &\quad \text{low} \end{aligned} \tag{33}$$

Equivalently, at both utilizations,

$$\gamma_i = 1 - \sum_{r<i} \phi_r \tag{34}$$

Note also that at low utilization, the probability is $1/(F\gamma_i)$ that a request for size i is met from the long free block. The essential feature of the long fragment (really infinite) is that its length does not alter when it is trimmed or extended by a finite amount. The low utilization analysis thus implicitly assumes an infinite store finitely loaded.

Two cases arise as with clustering: oversize fit ($r > i$) and exact fit ($r = i$). In the former case, with probability $b_i Q_{ir}$ for each i and r, $f_r \leftarrow f_r - 1$ and $f_{r-i} \leftarrow f_{r-i} + 1$. In the latter case, with probability $b_i Q_{ii}$ for each i, $f_i \leftarrow f_i - 1$.

The expected increment $\Delta^+\tilde{f}$ to $\tilde{f}$ following a single reservation is thus given by

$$\Delta^+\tilde{f} = \sum_i b_i \left\{ -Q_{ii}a^i + \sum_{r>i} Q_{ir}(a^{r-i} - a^r) \right\} \tag{35}$$

The probability p of an oversize fit is seen to be

$$p = 1 - \sum_i b_i Q_{ii} \tag{36}$$

2.2.2 Releases. Each reservation is equally likely to be released next and the probability is taken to be b_i that a block to be released is of size i. The numbers of reservations with 0,1,2 free neighbours are denoted by n_0, n_1, n_2. Evidently

$$n_0 + n_1 + n_2 = B \tag{37}$$

and

$$\tfrac{1}{2}n_1 + n_2 = F \tag{38}$$

Let p_0, p_1, p_2 be the probabilities that a reservation has 0,1,2 free neighbours. Then

$$p_0 = n_0/B, \quad p_1 = n_1/B, \quad p_2 = n_2/B \tag{39}$$

where

$$p_0 + p_1 + p_2 = 1 \tag{40}$$

and

$$p_1 + 2p_2 = x \tag{41}$$

Furthermore, referring back to the cluster analysis, it is clear that n_2 and c_1 are merely different notations for the count of lone reservations so that

$$p_2 = \frac{n_2}{B} = \frac{c_1}{F}\frac{F}{B} = \tfrac{1}{2}x\sigma_1 \tag{42}$$

and therefore, using Eqn (24),

$$p_2 = x^2/(x+4) \tag{43}$$

As with clustering, three cases arise.

(i) *No compaction.* Suppose the released block is an interior member of a cluster. Then with probability $b_i p_0$ for each i, $f_i \leftarrow f_i + 1$.

(ii) *Single compaction.* Suppose the released block is an end member of a cluster, and that the adjacent free

fragment is of size r. Then with probability $b_i p_1 \phi_r$ for each i and r, $f_r \leftarrow f_r - 1$ and $f_{r+i} \leftarrow f_{r+i} + 1$.

In the low utilization case, single compaction involving the long fragment does not change any of the block counts.

(iii) *Double compaction.* Suppose the released block is a lone block with left and right free neighbours of sizes r and s. Then with probability $b_i p_2 \Phi_{rs}$ for each i, r and s, $f_r \leftarrow f_r - 1, f_s \leftarrow f_s - 1$ and $f_{r+s+i} \leftarrow f_{r+s+i} + 1$ where

$$\Phi_{rs} = \frac{f_r}{F}\frac{f_s}{F-1} = \frac{F}{F-1}\phi_r\phi_s \quad \text{if } r \neq s$$
$$= \frac{f_r}{F}\frac{f_r - 1}{F-1} = \frac{F}{F-1}\phi_r\left(\phi_r - \frac{1}{F}\right) \quad \text{if } r = s \qquad (44)$$

In the low utilization case there is also the possibility of a double compaction involving the long fragment and one of length r. Thus with probability $2b_i p_2 \phi_r/(F-1)$, $f_r \leftarrow f_r - 1$.

Combining these results, the expected increment $\Delta^- \tilde{f}$ to $\tilde{f}$ following a single release may be written down. First at high utilization,

$$\Delta^- \tilde{f} = \sum_i b_i \left\{ p_0 a^i + p_1 \sum_r \phi_r (a^{r+i} - a^r) + p_2 \sum_r \sum_s \Phi_{rs}(a^{r+s+i} - a^r - a^s) \right\} \qquad (45)$$

which reduces to

$$\Delta^- \tilde{f} = (p_0 + p_1\tilde{\phi} + p_2\tilde{\phi}^2)\tilde{b} - x\tilde{\phi} + \frac{p_2}{F-1}(\tilde{\phi}^2 - \tilde{\phi}(a^2))\tilde{b} \qquad (46)$$

Similarly at low utilization,

$$\Delta^- \tilde{f} = \sum_i b_i \left\{ p_0 a^i + p_1 \sum_r \phi_r (a^{r+i} - a^r) + p_2 \sum_r \sum_s \Phi_{rs}(a^{r+s+i} - a^r - a^s) - \frac{2p_2}{F-1}\sum_r \phi_r a^r \right\} \qquad (47)$$

which, making use of the different normalization of $\tilde{\phi}$ (Eqn (29)), reduces to the same form, Eqn (46), as at high utilization.

2.2.3 Equilibrium. The equilibrium condition follows the same pattern as for clustering. Thus

$$\Delta^+ \tilde{f} + \Delta^- \tilde{f} = 0 \qquad (48)$$

2.2.4 The fifty per cent rule. Setting $a = 1$ in the equilibrium condition recovers Eqn (19) at high utilization and gives

$$\Delta^+(F-1) + \Delta^-(F-1) = 0 \qquad (49)$$

at low utilization, in each case expressing the constancy of F. On substituting into the expressions for $\Delta^+ \tilde{f}$ and $\Delta^- \tilde{f}$, the fifty per cent rule, Eqn (20), is recovered at both utilizations.

2.2.5 Transformation of variables. Using the expressions for $\Delta^+ \tilde{f}$ and $\Delta^- \tilde{f}$ in the equilibrium condition produces a set of fragmentation equations which determine the size distribution $\tilde{\phi}$. Because of the terms γ_i in the expression for $\Delta^+ \tilde{f}$, these equations cannot be expressed cleanly in generating function form. There seems little prospect in general of deriving analytical solutions and so attention has been concentrated upon numerical methods.

There are formally an infinite number of equations in an infinite number of unknowns. Given an arbitrary small quantity ε, it is reasonable to truncate the ϕ sequence after ϕ_l, say, where

$$\sum_{r>l} \phi_r < \varepsilon \qquad (50)$$

for then the probability is at most ε that a fragment longer than l is neglected. The long fragment at low utilization is safeguarded since it is kept separate from the ϕ sequence.

Partly for this reason, and partly because of the experience of Part 1, it was decided to transform variables from ϕs to χs where

$$\chi_i = \sum_{r>i} \phi_r, \quad i \geq 0 \qquad (51)$$

The χs form a decreasing sequence of non-negative elements having limit zero and

$$\chi_0 = 1, \quad \text{at high utilization}$$
$$= 1 - 1/F, \quad \text{at low utilization} \qquad (52)$$

The truncation condition is now more conveniently

$$\chi_l < \varepsilon \qquad (53)$$

The transformation may be expressed in generating function form by defining

$$\tilde{\chi}(a) = \sum_{r \geq 0} \chi_r a^r \qquad (54)$$

so that

$$\tilde{\phi} = \chi_0 - (1-a)\tilde{\chi} \qquad (55)$$

and

$$\phi_r = \chi_{r-1} - \chi_r \qquad (56)$$

2.2.6 The fragmentation equations. Making the transformation and noting that

$$\gamma_i = \chi_{i-1}, \quad \text{high}$$
$$= \chi_{i-1} + 1/F, \quad \text{low} \qquad (57)$$

it follows, after some algebra, that for each utilization,

$$\Delta^+ \tilde{f} = \sum_i \frac{b_i}{\gamma_i}\left\{ -\chi_{i-1}a^i + \chi_i + (1-a)\sum_{r \geq i} \chi_r (a^r - a^{r-i}) \right\} \qquad (58)$$

and that at high utilization,

$$\Delta^- \tilde{f} = \tilde{b} - x - (1-a)\left[p_1\tilde{b} - x + \left(\frac{F}{F-1}\right)2p_2\tilde{b} \right]\tilde{\chi} + (1-a)^2\left(\frac{F}{F-1}\right)p_2\tilde{b}\tilde{\chi}^2 + \frac{1-a^2}{F-1}p_2\tilde{b}\tilde{\chi}(a^2) \qquad (59)$$

and at low utilization,

$$\Delta^- \tilde{f} = \left(1 - \frac{x}{F}\right)\tilde{b} - x\left(1 - \frac{1}{F}\right) + (1-a)(1-\tilde{b})x\tilde{\chi} + (1-a)^2\left(\frac{F}{F-1}\right)p_2\tilde{b}\tilde{\chi}^2 + \frac{1-a^2}{F-1}p_2\tilde{b}\tilde{\chi}(a^2) \qquad (60)$$

At each utilization level, the results are next collected to form the equilibrium condition. In the simplification that follows, the fifty per cent rule $x = p$ is employed, where in terms of $\tilde{\chi}$

$$p = \sum_i b_i \chi_i / \chi_{i-1}, \quad \text{high}$$
$$= \sum_i b_i (\chi_i + 1/F)/(\chi_{i-1} + 1/F), \quad \text{low} \qquad (61)$$

This permits a factor $(1 - a)$ to be removed and leads at high utilization to the form

$$\sum_i \frac{b_i}{\gamma_i} \sum_{r \geq i} \chi_r (a^r - a^{r-1}) + \left[(1 - \tilde{b})x - \frac{2p_2}{F-1} \tilde{b} \right] \tilde{\chi} + (1 - a)\left(\frac{F}{F-1}\right) p_2 \tilde{b} \tilde{\chi}^2 + \frac{1+a}{F-1} p_2 \tilde{b} \tilde{\chi}(a^2) = 0 \qquad (62)$$

and at low utilization to

$$-\sum_i \frac{b_i}{\gamma_i} \left\{ \frac{1}{F}(1 + a + a^2 + \cdots + a^{i-1}) + \sum_{r \geq i} \chi_r (a^{r-i} - a^r) \right\} + \frac{x}{F} \tilde{\beta} + (1 - \tilde{b}) x \tilde{\chi} + (1 - a)\left(\frac{F}{F-1}\right) p_2 \tilde{b} \tilde{\chi}^2 + \frac{1+a}{F-1} p_2 \tilde{b} \tilde{\chi}(a^2) = 0 \qquad (63)$$

where

$$(1 - a)\tilde{\beta} = 1 - \tilde{b} \qquad (64)$$

so that

$$\tilde{\beta} = \sum_{r \geq 0} \beta_r a^r \qquad (65)$$

and

$$\beta_i = \sum_{r > i} b_r \qquad (66)$$

These forms are now separated into their component equations by extracting the coefficient of a^n for $n \geq 1$. After further manipulation a set of equations common to the two levels of utilization is obtained:

$$\left(x + \sum_{i \leq n} \frac{b_i}{\gamma_i} \right) \chi_n = h_n + \sum_i b_i \left\{ c_1 \chi_{n-i} - c_2 \chi_{(n-i)/2} + c_3 (U_{n-i-1} - U_{n-i}) + \frac{\chi_{n+i}}{\gamma_i} \right\} \qquad (67)$$

where

$$h_n = 0, \quad \text{high}$$
$$= \frac{1}{F} \sum_{i > n} b_i \left(\frac{1}{\gamma_i} - x \right), \quad \text{low} \qquad (68)$$

$$c_1 = x + \frac{2p_2}{F-1}, \quad \text{high}$$
$$= x, \quad \text{low} \qquad (69)$$

$$c_2 = \frac{p_2}{F-1} \qquad (70)$$

$$c_3 = \left(\frac{F}{F-1} \right) p_2 \qquad (71)$$

$$U_n = \sum_r \chi_r \chi_{n-r} \qquad (72)$$

and where [ ] denotes the integer part.

3. METHODS OF SOLUTION

The numerical methods used to solve Eqns (23) for clustering and (67) for fragmentation at the two levels of utilization are now described.

3.1 The cluster equations

For each n, $n \geq 1$, Eqn (23) determines σ_n in terms of σ_r for $r < n$. The sequence of σ values is thus easily generated in order of increasing subscript. In principle, therefore, a solution may be generated for any pair (x, F), or equivalently (B, F). Since p is a probability, the fifty per cent rule determines that $0 \leq x \leq 1$. Thus, feasible parameters must satisfy $F \leq B/2$.

3.2 The fragmentation equations

For each n, $n \geq 1$, Eqn (67) involves variables χ_r for $r \leq n + R$ where R is the size of the largest request. There being no evident direct solution, an iterative method was adopted. The parameters of an equation set, noting the result (43) for p_2, are seen to be x and F.

Suppose that χ_r, $r \geq 0$, is an approximate solution with χ_0 satisfying Eqn (52) and with $\chi_r = 0$ for $r > l$. The proposed iteration computes an improved approximation χ_r^*, $r \geq 0$, where $\chi_0^* = \chi_0$ and $\chi_r^* = 0$ for $r > l^*$. Equation (67) is modified by substituting starred quantities for χ and γ on the left hand side. By working through in order of increasing n, the nth equation thus permits χ_n^* to be calculated in terms of the existing χ values and of already calculated χ_r^*, $r < n$.

It is apparent that this will produce a larger set of χ^* values. By considering the maximum n for which χ_n^* can receive non-zero contributions, it is seen that $\chi_r^* = 0$ for $r > 2l + R + 1$. In practice, subject to space being available χ_n^* is computed for $n \leq 2l + R + 1$ and then the tail is docked to determine l^* such that $|\chi_r^*| \geq \varepsilon$ for $r = l^*$ and $|\chi_r^*| < \varepsilon$ for $r > l^*$ where ε is a small quantity.

Iterations are continued until, for each n, $|\chi_n^* - \chi_n| < \varepsilon$. Convergence has not been proved but has been achieved in practice.

3.2.1 Acceleration of convergence. The limiting rate of convergence was found to decrease as ε was reduced and it was therefore worth introducing an occasional acceleration step into the iteration sequence.

Let $\tilde{\chi}^{(i)}$ denote the ith approximation to $\tilde{\chi}$ and write, for some q,

$$\tilde{\Delta}1 = \tilde{\chi}^{(i+q)} - \tilde{\chi}^{(i)} \qquad (73)$$

and

$$\tilde{\Delta}2 = \tilde{\chi}^{(i+2q)} - \tilde{\chi}^{(i+q)} \qquad (74)$$

In order to apply an acceleration, $\tilde{\Delta}1$ and $\tilde{\Delta}2$ must be sufficiently parallel. The condition adopted is

$$(\tilde{\Delta}1:\tilde{\Delta}2)^2 \geq 0.99(\tilde{\Delta}1:\tilde{\Delta}1)(\tilde{\Delta}2:\tilde{\Delta}2) \quad (75)$$

where

$$(\tilde{P}:\tilde{Q}) = \sum_{r\geq 1} P_r Q_r \quad (76)$$

This condition warrants writing, for $k = 0,1,2$,

$$\tilde{\chi}^{(i+kq)} = \tilde{A} + \lambda^k \tilde{H} \quad (77)$$

where $\tilde{A}$ is the required limiting value for $\tilde{\chi}$ and $\tilde{H}$ is the limiting direction of approach to the solution. If m denotes the subscript value of the dominant component of $\tilde{\Delta}2$, the rate factor λ may be estimated as

$$\lambda = \Delta 2_m/\Delta 1_m \quad (78)$$

and it follows that

$$\tilde{H} = \tilde{\Delta}2/(\lambda^2 - \lambda) \quad (79)$$

and therefore that

$$\tilde{A} = \tilde{\chi}^{(i+2q)} + \omega\tilde{\Delta}2 \quad (80)$$

where

$$\omega = \lambda/(1 - \lambda) \quad (81)$$

Precautions are necessary in applying an acceleration step. As the solution is approached, the $\tilde{\Delta}1$ and $\tilde{\Delta}2$ components get smaller and significant figures are lost, affecting the estimation of λ. The parameter q was introduced with this in mind. By grouping iterations into bundles of size q, $\tilde{\Delta}1$ and $\tilde{\Delta}2$ are smoothed and magnified. Satisfactory results were obtained with $q = 5$, 6 and 7 though no optimization was attempted.

Furthermore, the slow convergence when ε is small yields values of λ only slightly below unity. This causes ω to be both large and imprecise so that $\tilde{A}$ is not well determined. To mitigate these effects, if $\lambda > 0.9$, λ is replaced by 0.9.

A fuller discussion of this many-variable extension of the well-known Aitken acceleration technique is given in Ref. 5.

3.2.2 Low utilization. On the basis of the earlier exploratory work with unit length requests, it was expected that at high utilizations satisfactory solutions would be found for sufficiently low x values with each choice of F. At low utilizations, the physically controllable parameter is B, the mean number of reservations. It was expected that for each B, the model would determine both F and x.

On carrying out the iterative solution of the equations with given B and trial F, and hence determining p from Eqn (61), it was observed that according as F is chosen too small or too large, so p is greater than or less than x. It seemed therefore that a simple binary chop outer iteration would serve to refine the interval $1 \leq F \leq B/2$ within which F must be located to satisfy the fifty per cent rule $x = p$.

In the y–F plane, the curve $y = x$ is the straight line $y = (2/B)F$. Near the intersection $x = p$, the curve $y = p$ is also approximately linear, its slope being only a little less than $(2/B)$. A small error in the determination of the p curve can, therefore, significantly affect the location of its intersection with the x line.

Each evaluation of p is subject to appreciable error and so binary chop is not an appropriate technique. Rather, it was decided to compute the least squares linear approximation to the p curve from evaluations of p at a range of points near where $x - p$ changes sign. By this means computational errors would be contained rather than further magnified.

The limiting value of x as B increases should identify the threshold separating the high and low utilization regimes.

4. CALCULATIONS OF CLUSTERING

The case $x = 1$ is special in that Eqn (21) reduces to

$$a(1 - a)\tilde{\sigma}' + 2(2 - 2a + a^2)\tilde{\sigma} - 2a = 0 \quad (82)$$

which is independent of F. As noted in Part 2, this equation has the analytic solution

$$\tilde{\sigma} = \frac{1}{2a^4}\{3(1 - a)^2 e^{2a} + 2a^3 + 3a^2 - 3\} \quad (83)$$

leading to

$$\sigma_n = 3 \times 2^{n+1}\frac{n(n + 3)}{(n + 4)!} \quad (84)$$

Calculations based upon Section 3.1 were carried out for a range of x and F values, the special case above providing a simple check. A striking feature of the results is their insensitivity to F over a wide range. Thus the validity of the model derived in Part 2 for infinite F is extended effectively to finite F. In particular the two so-called geometric approximations

$$\tilde{\sigma} \approx (1 - \eta)a/(1 - a\eta) \quad (85)$$

where we may either choose η to match the theoretical mean sequence length, giving

$$\eta = 1 - x/2 \quad (86)$$

or, possibly more importantly, to preserve the proportion of lone reservations, giving

$$\eta = 1 - 2x/(x + 4) \quad (87)$$

The calculations also indicate the range of validity of the present model. At low F values the solution $\tilde{\sigma}$ is inadmissible since the σ_n values no longer rise monotonically to 1.0 but overshoot and then execute damped oscillations about 1.0 as n increases further. The critical F value drops from about 60 at $x = 0.2$ to 30 at $x = 0.9$. The σ_n values decay most rapidly with n at high x and so the observed behaviour is consistent with a need to maintain some minimum level of the values of the significant c_n counts at each x value.

A possible explanation is that the assumption of perfect mixing at each reservation and release, implicit for example in the use of ψ_{rs} in Section 2.1, is suspect. At low c-counts the behaviour is dependent upon the recent history. The appropriate model at low occupation numbers would be a Markov process where different spatial configurations of clusters and fragments are explicitly recognized as separate states with individual transition probabilities.[4] These considerations apply of course with equal force to the analysis of fragmentation.

5. VALIDATION USING UNIT LENGTH REQUESTS

Fragmentation was calculated for unit length requests using both direct and iterative solution methods. The primary interest was to assess the accuracy of the iterative method as a guide for use when applying it to general request distributions.

Computing was carried out on the ICL 1906S at Liverpool, with eleven significant decimal digit precision. Differencing checks indicated that χ components were evaluated in the iteration loop with, typically, an accuracy of nine decimal places. With $\varepsilon = 0.5 \times 10^{-6}$ convergence was prohibitively slow at low utilization without acceleration and the noise produced by an acceleration was comparable to ε in magnitude. Thus the solution process is effectively limited to an accuracy of about five decimal places in the low utilization region. At high utilization the tail of χ is shorter and greater accuracy is obtained.

5.1 A direct solution

As discussed in Parts 1 and 2, the case of unit length requests is special in that a direct solution of the fragmentation equations is possible. Substitution of $\tilde{b} = a$ into Eqn (67) yields, for $n \geq 1$,

$$(x+1)\chi_n = c_1\chi_{n-1} - c_2\chi_{[(n-1)/2]} + c_3(U_{n-2} - U_{n-1}) + \chi_{n+1} \qquad (88)$$

where c_1, c_2 and c_3 are as defined in Eqns (69)–(71). In addition, the fifty per cent rule and the known form for χ_0 combine to determine χ_0 and χ_1. Thus at high utilizations,

$$\chi_0 = 1, \qquad \chi_1 = x \qquad (89)$$

and at low utilizations

$$\chi_0 = 1 - 1/F, \qquad \chi_1 = x - 1/F \qquad (90)$$

Therefore, for $n = 1, 2, 3, \ldots$ in turn, Eqn (68) enables χ_{n+1} to be evaluated in terms of elements already available.

The estimation of F for each B at low utilization was discussed in Part 2, section 7.5. The analysis has already enforced the fifty per cent rule and a criterion based upon the form of the χ sequence is used instead. For given B and trial F, the generation of the χ sequence is terminated if either $\chi_{n+1} < 0$ or $\chi_{n+1} > \chi_n$. It is observed that these correspond, respectively, to F being too small or too large. In order to provide maximum contrast with the general program, and because of the greater accuracy expected of a direct calculation, binary chop was judged suitable as a means of refining F for each B.

The program for the direct calculation was itself checked over a range of x values at high utilization with $F = \infty$ against the analytical expressions for the χ elements given in Part 1, Eqns (3.37) and (3.39). Listings were produced with a precision of six decimal places and complete agreement was observed.

5.2 Numerical results

At high utilization, iterative solutions were obtained for comparison with direct solutions at various x values and at $F = \infty$ and $F = 128$. For each combination, calculations were made with both four and six decimal places nominal precision, i.e. $\varepsilon = 0.000\,05$ and $\varepsilon = 000\,0005$. At $x = 0.1$ the agreement was virtually exact. At $x = 0.4$ the iterative solution lost a decimal place and at $x = 0.5$ it lost two places in terms of the nominal precision. Thus at $x = 0.5$ the iterated solution with $\varepsilon = 0.000\,000\,5$ was accurate to four decimal places in χ. As x increases to its limiting value 0.5151 the performance deteriorates further. The variation of χ with F is negligible at $x = 0.1$ and rises to about 0.001 at $x = 0.5$.

At low utilization comparisons were made for various B values. Calculations were made with both four and five decimal places nominal precision. At $B = 64$ the iterated solution loses about two decimal places in χ, rising to three at $B = 512$. Thus at $B = 512$ the iterated solution with $\varepsilon = 0.000\,005$ was accurate to two decimal places only in χ. Above $B = 512$ the performance deteriorates further. A plot of x against $1/B$, however, permits the limiting value of x as $B \to \infty$ to be extrapolated—see the curve $R = 1$ in Fig. 2.

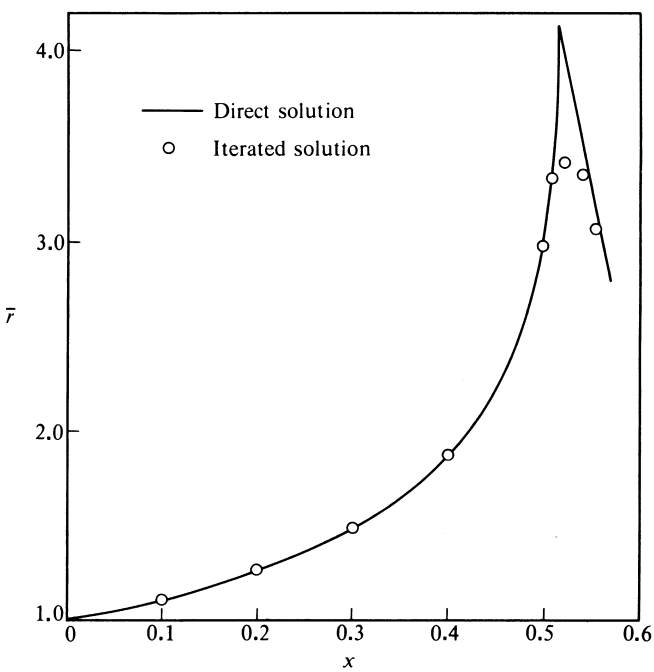

Figure 1. Validation with unit length requests.

The mean size $\bar{r}$ of the finite fragments is plotted against x in Fig. 1. The full curve represents results from the direct method and shows the peak in $\bar{r}$ at $x = 0.5151$ separating the high and low utilization regions. Also plotted are individual results from the iterative method, corresponding to the higher nominal precision in each region. Clearly the size of the peak in $\bar{r}$ is not well determined but its location in x is relatively more accurately established by the iterations at the two levels of utilization. The steepness of the right hand face indicates that at low utilizations x is relatively insensitive to B.

6. THREE SIMPLE EXPERIMENTS

The successful validation of the model enables it to be employed as a predictive tool. Three simple experiments

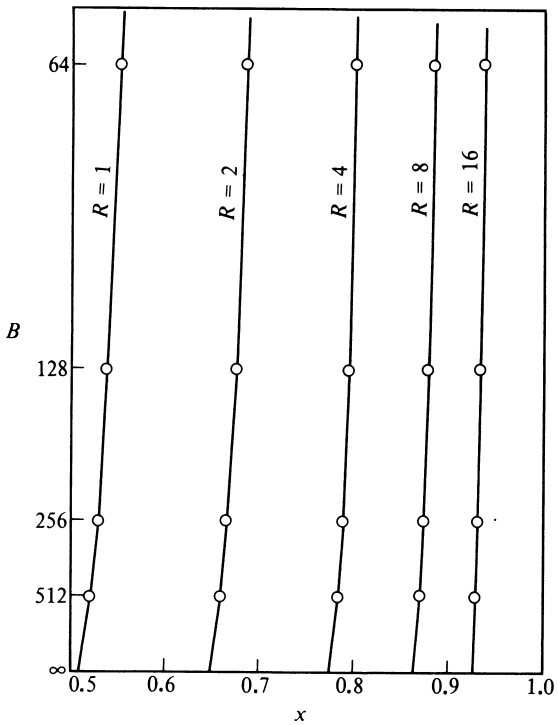

Figure 2. x vs B for uniform requests.

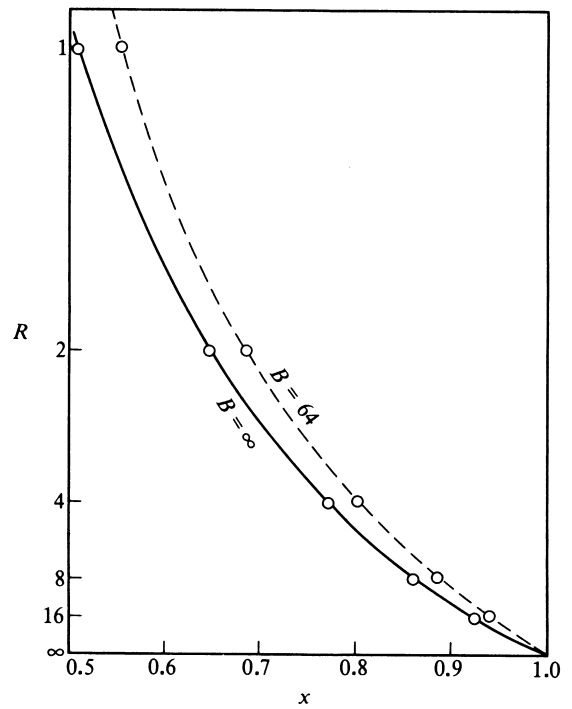

Figure 3. x vs R for uniform requests.

are now described which explore the influence of request-size distribution upon fragmentation.

6.1 Uniform requests

A sequence of calculations was carried out for the case of uniformly distributed requests, for which $b_r = 1/R$, $1 \leq r \leq R$. The cases $R = 1$, 2, 4, 8, 16 were treated. Of particular interest was the variation with R of the threshold x-value separating the high and low utilization regimes. For each R, low utilization calculations were made for $B = 64$, 128, 256, 512. These values lie within the region of validity of the model.

Table 1. Variation of x with R and B

R	B				
	64	128	256	512	∞
1	0.555	0.540	0.530	0.522	0.51
2	0.688	0.677	0.668	0.660	0.65
4	0.804	0.797	0.790	0.786	0.77
8	0.887	0.880	0.875	0.870	0.86
16	0.939	0.936	0.933	0.929	0.92

The computed x-values are shown in Table 1 and are plotted in Fig. 2 in the $x{:}B^{-1}$ plane to facilitate extrapolation to $B = \infty$ at each R value. These are the required x thresholds and form the right-hand column of the table. Figure 3 shows similarly in the $x{:}R^{-1}$ plane the plots for $B = 64$ and $B = \infty$ to permit extrapolation to $R = \infty$. The two curves lie close together and the other B curves, not shown, would lie between them. It is seen that, for each B, the limiting value of x as $R \to \infty$ is 1. This confirms Knuth's conjecture[3] when formulating the fifty per cent rule. The present work adds to that conjecture an indication of the rate of approach to the limit as the variability of request sizes increases.

Also of interest is the variation with R of the threshold value of the utilization, θ. Calculations were made at high utilization and $F = \infty$ for a range of x values. Computed values of θ are plotted against x for each R in Fig. 4. As the threshold is approached, the tail of χ lengthens and the values obtained are less secure. In the simple case $R = 1$, of unit length requests, the earlier papers showed that analytic solutions are available for the χ-equations. These locate the threshold at $x = 0.5151$, $\theta = 0.4849$ and it is relatively simple to show that, at this threshold, $dx/d\theta = 0$. For other R values the threshold x values were obtained from the extrapolations of Fig. 2 and it is tempting to conjecture that here too the tangents are parallel to the θ axis. The analysis is more complex and no formal verification has been obtained.

It will be appreciated that the extrapolated threshold θ values are subject to numerical uncertainty. This is deliberately highlighted in Fig. 4 by leaving the loose

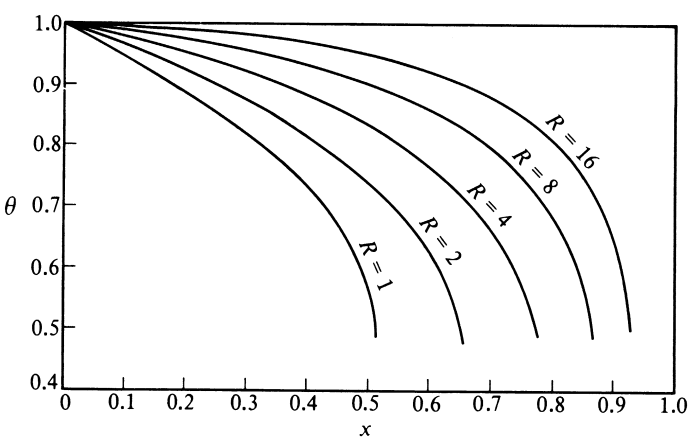

Figure 4. θ vs x for uniform requests.

ends of the curves somewhat ragged. What does emerge quite clearly however is the relative insensitivity of the threshold utilization to the request distribution. This is confirmed by the low utilization results. All estimates of the critical utilization lie between 45% and 50%.

6.2 Bipolar requests

The second experiment investigates, for given mean request size, the effect upon fragmentation of changes in the variance of the request distribution. For simplicity, the cases studied have only two different sizes of request, each occurring with the same probability. Thus, choosing 8 as the mean size and denoting the standard deviation by σ, the request distribution is specified by $b_{8+\sigma} = b_{8-\sigma} = 0.5$ with all other elements zero. Calculations were carried out for $0 \leq \sigma \leq 7$.

With certain request distributions, and on the assumption that the store pattern is consistent with an initial priming from an empty store, certain fragment sizes will never arise. For example, if all requests are of even length, there will be no free fragments of odd length. The choice of initial configuration when solving the fragmentation equations should bear this in mind since the iterations will not remove incompatible sizes.

It is readily seen that achievable sizes of fragment are of the form $\sum r n_r$ where the sum is over the values r for which $b_r \neq 0$ and where the n_r are arbitrary signed integers. It follows that if the request sizes have a common factor then all fragment sizes also have this factor, and that if the request sizes have no common factor then there is no restriction on fragment size.

In practice it is found that the presence of a common factor upsets the conditioning of the iterations. It is desirable, and clearly more efficient, to remove any common factor by rescaling the unit of length. Thus for example the case $\sigma = 4$ with $R = 12$ above is converted into $b_1 = b_3 = 0.5$ with $R = 3$, where each unit is four of the original words.

The results reproduced many of the features exemplified in Figs 2 and 4. Thus, at high utilization, the utilization θ drops from 1.0 at $x = 0$ to a value close to 0.5 at the threshold x where, as near as can be determined, the tangent is again parallel to the θ axis. At low utilization, as in Fig. 2, x decreases very slowly as B increases and approaches its threshold value as $B \to \infty$.

Table 2 lists the threshold x and θ for each σ. θ is not well determined by the calculations and the variation is probably not significant. The interesting features of the x threshold however are that it reflects the factoring of the request sizes and is largely insensitive to σ within each factor class. The first effect is readily explicable in terms of the filtering of achievable block sizes consequent upon a common factor, but the near constancy within each class is unexpected.

Table 2. Variation of threshold with σ

σ	0	1	2	3	4	5	6	7
x	0.52	0.79	0.76	0.79	0.72	0.80	0.78	0.80
θ	0.48	0.52	0.52	0.50	0.45	0.49	0.47	0.48

6.3 Skew requests

In the third experiment, the relative probabilities were varied for a fixed set of request sizes. A simple example was chosen, based on the case $\sigma = 2$ of the previous section. Thus the request distribution is specified by $b_3 = \beta, b_5 = 1 - \beta$. Calculations were carried out for a number of values of β in the range $0 \leq \beta \leq 1$. Those for $\beta = 0.5$ reproduce the earlier set with $\sigma = 2$ and those for $\beta = 0$ and $\beta = 1$ are simply rescaled versions of unit length requests.

The results follow a similar pattern to those of the earlier experiments with the threshold θ value remaining close to 50%. The x threshold varies only slowly in the middle of the β range but dips sharply at the extremes. Figure 5 shows this variation, extrapolated with the low

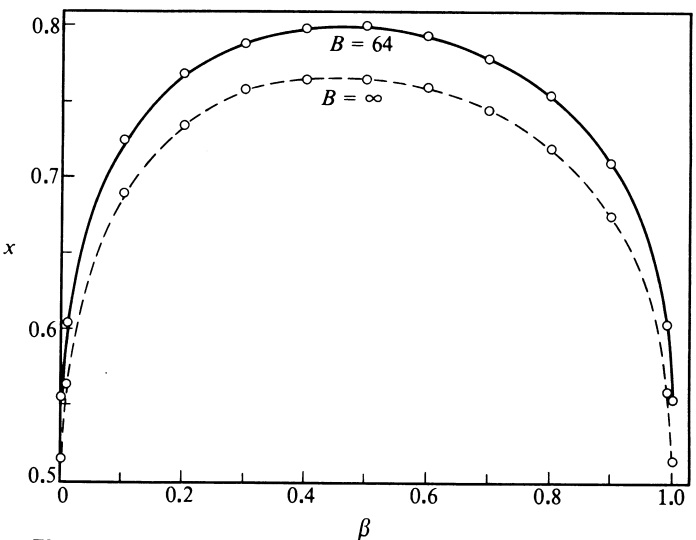

Figure 5.

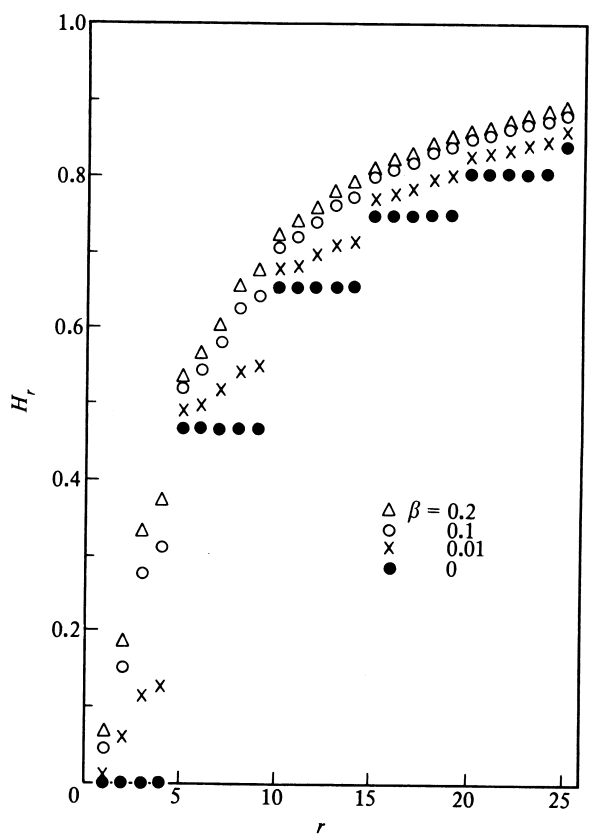

Figure 6. H_r vs r for skew requests at $B = 256$.

utilization model from calculations as before at $B = 64$, 128, 256, and 512. The values of x for $B = 64$ are also plotted to indicate the relative insensitivity of x to variations in B at each β. The interesting observation is that x appears to approach its limiting values continuously at the two ends of the β range.

In all three experiments, location of the utilization threshold has been given special consideration. This provides a severe test of the consistency of the models since the conditioning of the equations is worst near the threshold. A more general comparison is made in Fig. 6 by plotting the cumulative distribution of fragment sizes obtained for $B = 256$ from the low utilization model near $\beta = 0$. This shows quite clearly, at the larger fragment sizes, how the comparatively smooth curve at $\beta = 0.2$ evolves into a staircase function at $\beta = 0$. The sizes tend to favour multiples of 5 for each β, with the influence of 3s clearly discernible as β increases.

Acknowledgements

During the early stages of the work reported here, I made short visits to the University of Newcastle upon Tyne where I was able to discuss with Dr Terry Betteridge our different approaches to the analysis of store fragmentation. I am grateful to him and to Professor Brian Randall, Professor Harry Whitfield and their colleagues for their interest, conversation and hospitality.

REFERENCES

1. C. M. Reeves, Free store distribution under random fit allocation, Part 1. *The Computer Journal* **22**, 346–351 (1979).
2. C. M. Reeves, Free store distribution under random fit allocation, Part 2. *The Computer Journal* **23**, 298–306 (1980).
3. D. E. Knuth, *The Art of Computer Programming*, Vol 1: *Fundamental Algorithms*, 1st Edn, §2.5 Dynamic storage allocation, pp. 435–455 (1968).
4. T. Betteridge, An algebraic analysis of storage fragmentation. PhD Thesis, University of Newcastle upon Tyne (1979).
5. C. M. Reeves, A note on the acceleration of convergence. *IUCC Bulletin* **2**, 122–124 (1980).

Received January 1982