M. C. Easton

Model for Interactive Data Base Reference String

Abstract: A particularly simple Markov chain model for a reference string is described. The model, which is only slightly more complicated than the independent reference model, generates strings that have a locality property and that have a specified probability distribution of references over pages. Expressions are obtained for expected working-set size and expected working-set miss ratio. The model is used in an examination of the effect of grouping pages into blocks and in a discussion of the problem of evaluating the effect of changes in the size of the data base. Predictions of the model are shown to agree closely with observations of a string of data base references generated by an interactive data base system having a large number of concurrent users.

Introduction

In a two-level storage hierarchy, a reference to a page not in first-level storage is called a *miss*. The miss-ratio function (average number of misses per reference as a function of first-level storage capacity) is a basic tool for studying the performance of hierarchical storage systems [1–5]. The miss-ratio function depends on the replacement algorithm and on the properties of the reference string. In order to have some understanding of how variations in the replacement algorithm and/or the reference string affect the miss ratio, a model for the reference string is essential.

We examine here the structure of a reference string generated by an interactive data base system. The particular system studied, known as the Advanced Administrative System (AAS), has been described in some detail [6]. It is transaction-driven, and several hundred terminals are simultaneously active. The recorded reference string, which contains only the references to the data base (logical addresses of 1693-byte pages), is a composite of the activities at the various terminals. Thus, the string is qualitatively and quantitatively different from the single-thread program address traces used in most of the studies cited. For example, more than 10^5 distinct pages were referenced in the sample string. Program address traces typically contain at most several hundred distinct page references.

One motivation for studying references to the data base is the availability of on-line storage devices with extremely large capacities. (For example, devices exist that can store more than 10^{10} bytes of data.) Information concerning the structure of the data base reference string is a basic requirement for studies of performance of a system that uses such a device as a backing store for disk storage.

One of the simplest reference string models is the familiar independent reference model. At each discrete time point, a page reference is randomly chosen from the set of page addresses in accordance with a fixed probability distribution. Unfortunately, this model does not adequately reproduce the miss ratios of the sample AAS string. In particular, for AAS, the LRU replacement algorithm (which replaces the least-recently-used page) yields a smaller miss ratio than does the A_0 algorithm (which replaces the page having the least relative frequency of reference). These results are shown in Fig. 1. Under the independence assumption, the LRU miss ratio cannot be smaller than that using A_0 [7, 8].

An examination of the data suggests that dependence is primarily a local phenomenon. The transaction taking place at a particular terminal tends to re-reference a page that has already been referenced, but after several transactions at that terminal the probability of referencing that page is independent of previous activity. This behavior, which is consistent with the "principle of locality" [9], was modeled by a particularly simple Markov chain. Because of the simplicity, analytic results are obtainable for expected working-set sizes and miss ratios, and these closely approximate the observed values if the window size is reasonably large. (Large window sizes are of interest because of the application to systems having 10^6 or more bytes of data in first-level storage.) The model permits consideration of such problems as dependencies of miss ratio on page size and on size of the data base.

Before describing the model, we give some basic definitions and assumptions and a few useful relationships.

Working-set properties

We assume that the records in the data base are arranged in equal-sized *pages*. Each reference to a record is associated with the address of the page that contains it. Let $X = \{x_1, x_2, \cdots, x_n\}$ be the set of page addresses.

We define a *reference sequence* to be a sequence of random variables $\{R_1, R_2, \cdots, R_t, \cdots\}$ each having as range the set X. A *reference string* is a realization of a reference sequence. We limit consideration here to *stationary sequences*. That is, for any positive integer m, any $\{R_{i_1}, \cdots, R_{i_m}\}$ and any $\{x_{j_1}, \cdots, x_{j_m}\}$, we have for any positive integer k

$$Pr\{R_{i_1} = x_{j_1}, \cdots, R_{i_m} = x_{j_m}\}$$
$$= \Pr\{R_{i_1+k} = x_{j_1}, \cdots, R_{i_m+k} = x_{j_m}\}.$$

Let T, the *window size*, be a positive integer. For each $t \geq T$, we define the random variable

$$M(t, T) = \begin{cases} 1 \text{ if the value of } R_{t+1} \text{ is distinct from} \\ \quad \text{those of } R_{t-T+1}, R_{t-T+2}, \cdots, \text{and } R_t, \\ 0 \text{ otherwise.} \end{cases}$$

The stationarity of $\{R_t\}$ implies the stationarity of $\{M(T, T), M(T+1, T), \cdots\}$, so the expected value of $M(t, T)$ is not a function of t. Define $\overline{M}(T) = E[M(t, T)]$; $\overline{M}(T)$ is called the *expected working-set miss ratio*, representing the probability that the value of R_{t+1} is distinct from the values of the T previous sequence members.

For $t \geq T$, let $\{S(t, T)\}$ be random variables defined as follows: $S(t, T)$ is the number of *distinct* values taken by the set $\{R_{t-T+1}, R_{t-T+2}, \cdots, R_t\}$. (Here $S(t, T)$ is the familiar working-set size with window size T, terminated at t.)

The stationarity of $\{S(T, T), S(T+1, T), \cdots\}$ follows from that of $\{R_t\}$. The *expected working-set size*, denoted by $\overline{S}(T)$, is $E[S(t, T)]$. We note that

$$S(t, T+1) = S(t, T) + M(t, T). \tag{1}$$

This imples that

$$\overline{S}(T+1) = \overline{S}(T) + \overline{M}(T). \tag{2}$$

Because $\overline{S}(1) = 1$, we have

$$\overline{S}(T) = 1 + \sum_{j=1}^{T-1} \overline{M}(j). \tag{3}$$

(Relations (2) and (3) were derived, using somewhat different methods, by Denning and Schwartz [9].)

We discuss briefly the question of estimating values of $\overline{S}(T)$ and $\overline{M}(T)$ from a reference string of length $j > T$. (For a detailed discussion, including efficient computational methods, see [10].) Define random variables

$$M_j(T) = \frac{1}{j-T} \sum_{t=T}^{j-1} M(t, T),$$

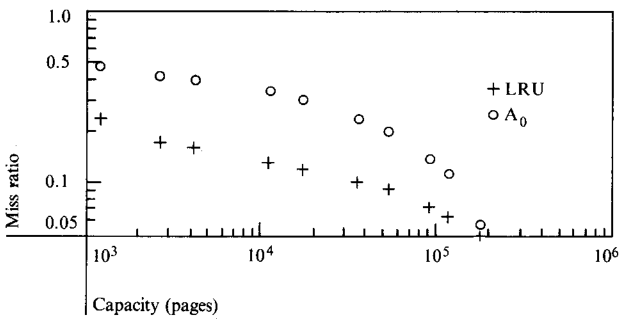

Figure 1 Comparison of LRU and A_0 miss ratios.

$$S_j(T) = \frac{1}{j-T+1} \sum_{t=T}^{j} S(t, T).$$

Note that the value of $M_j(T)$ is an average of values of $\{M(t, T)\}$ and that of $S_j(T)$ is an average of values of $\{S(t, T)\}$. By linearity of expectation, it follows that $M_j(T)$ is an unbiased estimator for $\overline{M}(T)$ and $S_j(T)$ is an unbiased estimator for $\overline{S}(T)$. (That is, $E[M_j(T)] = \overline{M}(T)$, $E[S_j(T)] = \overline{S}(T)$.) If we assume certain asymptotic correlation properties, we can also show that $S_j(T)$ converges in probability to $\overline{S}(T)$ and $M_j(T)$ converges in probability to $\overline{M}(T)$. These results follow from the discussion of the ergodic theorem in [11], with the assumption that for fixed T,

$$\lim_{x\to\infty} E[M(t, T)\ M(t+x, T)] = [\overline{M}(T)]^2 \text{ and}$$
$$\lim_{x\to\infty} E[S(t, T)\ S(t+x, T)] = [\overline{S}(T)]^2.$$

Previous work on modeling reference strings

As indicated in the introduction, most previous studies of properties of and models for reference strings have been concerned with single-thread program address traces. The characteristics of these strings are somewhat different from those of a string of references to the data base generated by an interactive system. Nevertheless, awareness of some aspects of success and failure in modeling program address traces is useful in considering the present work. Perhaps the most popular models in the literature are the independent reference model and the LRU stack model (see [12, ch. 6] for descriptions). For the programs studied in [13], the independent reference model does not adequately reproduce miss ratios. On the other hand, the LRU stack model, which has the (observed or assumed) LRU miss ratios built in, reproduces working-set miss ratios quite well. However, the stack model predicts uniform distribution of references over pages [12, ch. 6], which is contrary to observations of actual strings. In contrast, the model described here has the (observed or assumed) probability distribution of references over pages built in (as 551

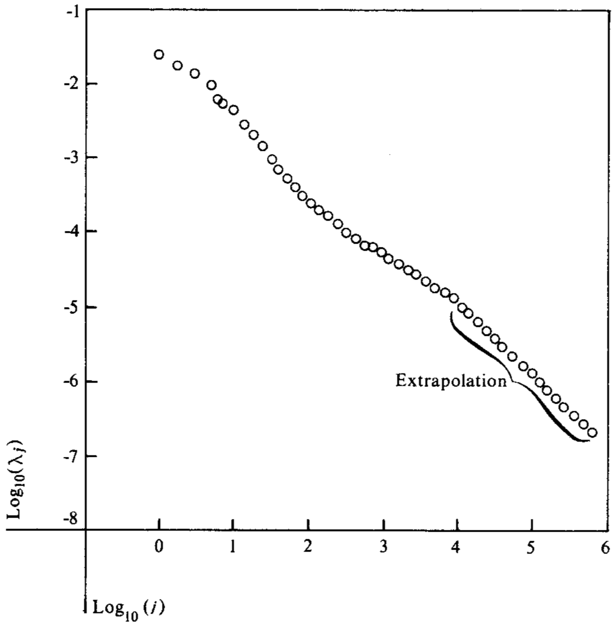

Figure 2 Observed reference frequency of jth most frequently referenced page.

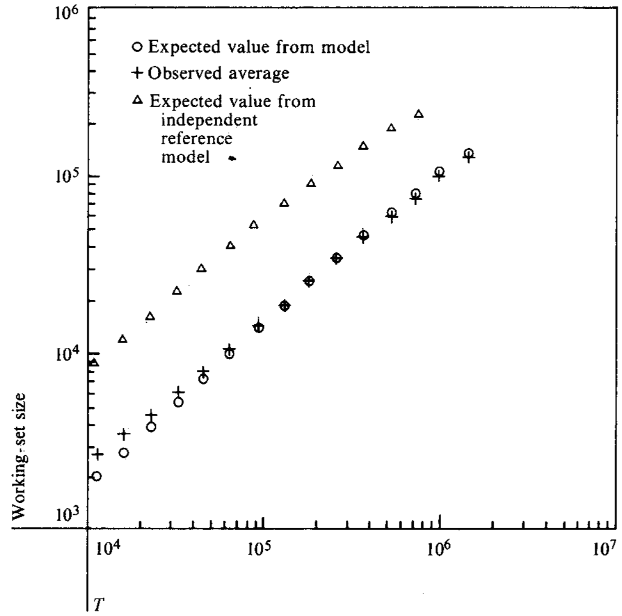

Figure 3 Comparison of observed average working-set sizes with expected values from models.

does the independent reference model) but also yields, through the structure of the model and one parameter, miss ratios close to observed values.

Model

It has been noted informally that observed reference strings often reference sequentially a set of physically adjacent (e.g., consecutively stored) records. If the records are stored in pages, this is generally seen as a sequence of references to the same page. On the other hand, a random access pattern is also frequently seen. The model described here attempts to capture both features.

The model is a finite Markov chain having n states, one for each member of X. Thus a sequence of states in the model corresponds to a reference string.

The transition probabilities are

$$p_{ii} = r + (1-r)\lambda_i,$$

$$p_{ij} = (1-r)\lambda_j,\ i \neq j, \tag{4}$$

where $0 \leq r < 1$, $\Sigma_{i=1}^{n} \lambda_i = 1$, and $\lambda_i > 0$ for $i = 1, \cdots, n$.

We can view the transition probabilities as arising from the following mechanism. With probability r, the previously referenced page is re-referenced. With probability $1 - r$, the next reference is chosen in accordance with probability distribution $\{\lambda_i\}$. The case $r = 0$ corresponds to the independent reference model; a large value of r indicates high probability of sequential (or local) references to records. (A generalized, but relatively intractable, form of this model appears in [14] and [15].)

By using the terminology of [16] and noting that $p_{ij} > 0$ for all i, j, it may be verified that the chain is irreducible and that all states are aperiodic and positive recurrent. Therefore, by Theorem 2.5 of [16], the chain has a unique stationary distribution determined by

$$z_j = \sum_{i=1}^{n} z_i\, p_{ij}.$$

In fact, it is easy to check that $z_i = \lambda_i$ is a (and hence the) solution. Thus the vector $(\lambda_1, \cdots, \lambda_n)$ is the stationary distribution of the chain. We define the reference sequence as follows: R_1 takes a value in accordance with the stationary distribution; R_{i+1} depends on R_i as indicated in (4). Then by I.2.1 of [16] the sequence is stationary. Also, it can be shown that $\{M(t, T)\}$ and $\{S(t, T)\}$ have the asymptotic correlation properties discussed in the previous section. (A related question is discussed in [17].)

We now obtain expressions for $\overline{S}(T)$ and $\overline{M}(T)$. Let $\{W_i(1, T), \cdots, W_i(t, T) \cdots\}$ be random variables defined by

$$W_i(t, T) = \begin{cases} 1, \text{ if } R_j = x_i \text{ for some } j = t - T + 1, \cdots, t, \\ 0, \text{ otherwise.} \end{cases}$$

Then

$$S(t, T) = \sum_{i=1}^{n} W_i(t, T).$$

By using the Markov property, we have

$$\begin{aligned} \mathrm{E}[W_i(t, T)] &= 1 - \Pr\{R_j \neq x_i \text{ for } j = t - T + 1, \cdots, t\} \\ &= 1 - \Pr\{R_{t-T+1} \neq x_i\} \\ &\quad \times \Pr\{R_{t-T+2} \neq x_i | R_{t-T+1} \neq x_i\} \cdots \\ &\quad \times \Pr\{R_t \neq x_i | R_{t-1} \neq x_i\} \\ &= 1 - (1 - \lambda_i)[1 - \lambda_i(1 - r)]^{T-1}. \end{aligned}$$

Thus

$$\overline{S}(T) = n - \sum_{i=1}^{n} (1 - \lambda_i)[1 - \lambda_i(1 - r)]^{T-1}, \tag{5}$$

and, by (2),

$$\overline{M}(T) = (1 - r) \sum_{i=1}^{n} \lambda_i(1 - \lambda_i)[1 - \lambda_i(1 - r)]^{T-1}. \tag{6}$$

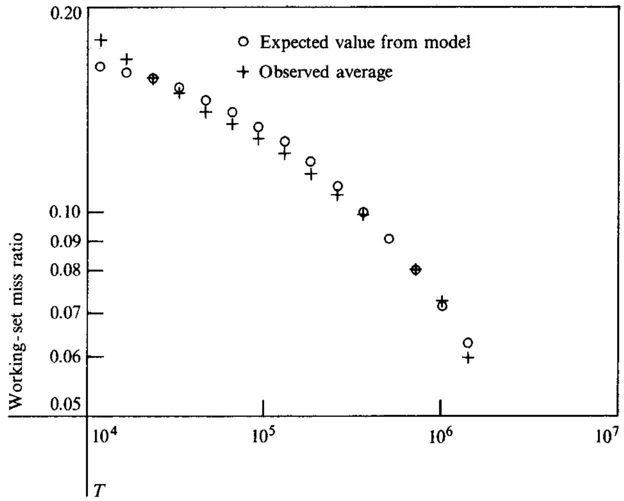

Figure 4 Comparison of observed average working-set miss ratios with expected values from model.

Comparison with AAS reference string

It would be quite remarkable if a model as simple as that given above could describe the detailed behavior of a data base reference string. The observations that led to the model, as outlined in the introduction, suggested that in the actual string a page reference has a high probability of being re-referenced within the next several transactions. Moreover, transactions are intermixed in the recorded string. In the model, the re-referencing phenomenon is represented by a high probability that two *successive* references are identical. Thus the model is not suited for comparison with measurements made for small values of T. A sequence of identical page references that the model generates successively would be expected to occur in an actual string spread out over some number, $\tilde{T}$, of references. (As indicated below, $\tilde{T} \approx 5 \times 10^4$.) It is reasonable to hope that the model represents the average contents of a string of length $\tilde{T}$ or larger.

As a result of the characteristics noted above, one cannot apply the simple procedure of estimating the value of r by measuring the relative frequency of immediate re-references, $1 - M_j(1)$, and applying (6) with $T = 1$. The same basic idea, however, can be used if we choose some $T_1 \geq \tilde{T}$ and apply (6) to the observed value of $M_j(T_1)$. In conjunction with the values of $\{\lambda_i\}$ described below, this was the method used to estimate the value of r.

The values of $\lambda_1 \geq \lambda_2 \geq \cdots$ were estimated from the observed relative frequencies. For the smaller relative frequencies, where one hundred or fewer references to a page appeared in a sample string, a power law extrapolation was used:

$$\lambda_j = c/j^{\alpha}. \tag{7}$$

The form of (7), with $\alpha \approx 1$, is suggested by several studies cited in [18] in a discussion of Zipf's law and the related "80-20" law (80 percent of data base references are to 20 percent of the records). Figure 2 shows the values of $\{\lambda_i\}$ used in the application of (5) and (6). At the time of the measurement, many records in the data base were obsolete, so it was not possible to determine accurately the number of pages. Thus the value of n was determined from the normalization condition $\Sigma_{i=1}^{n} \lambda_1 = 1$.

Figure 3 shows the values of $\overline{S}(T)$ computed from the model using $\alpha = 0.96$, $c = 0.08353$, $r = 0.79$, and values of $S_j(T)$ observed for $j = 2 \times 10^6$. Values of $\overline{S}(T)$ computed using $r = 0$ (independent reference model) are also shown. Figure 4 shows expected working-set miss ratios, $\overline{M}(T)$, computed using the same parameters, and observed values of $M_j(T)$ for the same reference string. (The values for the independent reference model are not plotted because most of the points would be off scale.)

To estimate the value of $\tilde{T}$ for AAS we note that, for the recorded string, the average number of references per transaction is about 40 and the average number of active terminals at any one time is about 400 [19]. Assuming re-referencing at one terminal has high probability for three transactions, we have $\tilde{T} \approx 4.8 \times 10^4$. Figures 3 and 4 show good agreement between model and observation for T larger than this number.

Extrapolating to larger data bases

It would be useful to have a method for projecting the performance of a data base system as the size of the data base grows. In this paragraph we assume for simplicity that the sequence $\{\lambda_j\}$ is adequately represented for *all*

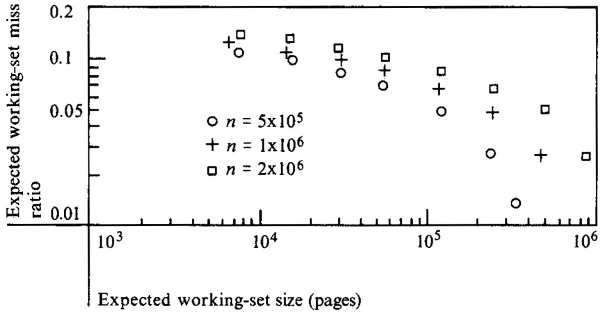

Figure 5 Expected working-set miss ratios for several data base sizes.

j by $\lambda_j = c/j^\alpha$. A reasonable assumption to make is that α above and parameter r of (4) will remain constant in time. Thus we need only project a future value of c, which is obtainable by estimating a future value of n and choosing c so the probabilities sum to one. For example, based on an estimate of the future workload of services performed by the system and some knowledge of the data structure it should be possible to estimate the future value of $\overline{S}(T_0)$ for one T_0. Then n can be chosen to bring the value predicted by the model close to $\overline{S}(T_0)$. (One should be aware of the sensitivity of the projection described above to the value of α assumed.)

Figure 5 shows expected working-set miss ratios for $\alpha = 0.9$, $r = 0.8$, and several values of n. In Fig. 5, the value of expected working-set miss ratio is plotted against corresponding expected working-set size in order to indicate the approximate capacity of first-level storage required to achieve a given miss ratio.

Partitioning pages into blocks

As indicated in the introduction, one aim of this study is to consider variations in the operation of the replacement algorithm. One possible alteration is as follows: For some integer k, called the *blocking factor*, partition the n pages into m blocks of k pages each. (If necessary, add dummy pages so that $mk = n$). The reference string is then viewed as a sequence of *block addresses* from a set $Y = \{y_1, \cdots, y_m\}$. For a fixed capacity of main storage, a replacement algorithm that operates on a sequence of block addresses and transfers data in blocks may exhibit a lower miss ratio than does the same algorithm operating on page addresses and transferring data in pages. We examine the following question: If expected working-set size is held constant (in pages), does working-set management of *blocks* results in lower miss ratios than does working-set management of *pages*? The effectiveness of this technique in reducing the miss ratio depends on properties of the reference string and on the partitioning method used.

We examine this question under the assumption that the reference string of page addresses is described by the model given previously. The reference string of block addresses is described by a Markov chain model having m states. We first find the transition probabilities.

For any choice of partition, we always index the pages so that the first block contains the pages with addresses $x_1, \cdots, x_k$, the second block contains the pages with addresses $x_{k+1}, \cdots, x_{2k}$, etc. Alternatively, the order chosen for the pages may be viewed as implicitly defining the partition. In either case, the new relative frequencies are

$$\lambda_j^* = \lambda_{(j-1)k+1} + \cdots + \lambda_{jk}, j = 1, \cdots, m.$$

Let p_{ij}^* be the probability that a reference to block j follows a reference to block i. Let $u \in \{1, \cdots, k\}$. Then

$$\begin{aligned}\Pr\{R_{t+1} = c \in \{x_1, \cdots, x_k\} | R_t = x_u\} &= \sum_{j=1}^{k} p_{uj} \\ &= (1-r)(\lambda_1^* - \lambda_u) + r + \lambda_u(1-r) \\ &= \lambda_1^*(1-r) + r.\end{aligned}$$

From this it follows that

$$p_{11}^* = r + \lambda_1^*(1-r).$$

Similarly,

$$\begin{aligned}&\Pr\{R_{t+1} = c \in \{x_{k+1}, \cdots, x_{2k}\} | R_t = x_u\} \\ &\quad = \sum_{j=k+1}^{2k} p_{uj} = (1-r)\lambda_2^*.\end{aligned}$$

Thus,

$$p_{12}^* = (1-r)\lambda_2^*.$$

Generalizing the above argument yields the following transition probabilities for the Markov chain:

$$p_{ii}^* = r + \lambda_i^*(1-r),$$

$$p_{ij}^* = (1-r)\lambda_j^*, i \neq j. \tag{8}$$

Also, we have $0 \leq r < 1$, $\sum_{i=1}^{m} \lambda_i^* = 1$, and $\lambda_i^* > 0$ for $i = 1, \cdots, m$.

The form is the same as (4) so (5) and (6) become

$$\overline{S}^*(T) = m - \sum_{i=1}^{m} (1-\lambda_i^*)[1-\lambda_i^*(1-r)]^{T-1}, \tag{9}$$

$$\overline{M}^*(T) = (1-r) \sum_{i=1}^{m} (\lambda_i^*)(1-\lambda_i^*)[1-\lambda_i^*(1-r)]^{T-1}, \tag{10}$$

where $\overline{S}^*(T)$ represents the expected number of distinct blocks in a string of length T and $\overline{M}^*(T)$ represents the probability that the current block reference has not occurred in the previous T references.

We have not yet considered the choice of method for partitioning the pages. For the case $r = 0$, this problem was studied by Yue and Wong [20]. In particular, they found a partition (described below) that minimizes $\overline{S}^*(T)$ for all T and gave several arguments for the use of this criterion. We show next that their partition minimizes $\overline{S}^*(T)$ for all T and all $1 > r \geq 0$.

We first note that the function $\phi(\lambda_i^*) = (1 - \lambda_i^*)[1 - \lambda_i^*(1-r)]^{T-1}$ is continuous and convex for $\lambda_i^* \geq 0$. It then follows from [21, section 3.17] that indexing the pages so that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ minimizes $\overline{S}^*(T)$ for all T. This means that pages are blocked together in nonincreasing order of probability of reference.

In order to compare miss ratios for two different values of k, we would like to choose window sizes so that the expected working-set sizes (measured in pages) are the same in both cases. If C is the expected working-set size in pages for the original case discussed ($k = 1$), then if $\overline{S}(T) = C$, we can write the expected miss ratio as $\overline{M}[\overline{S}^{-1}(C)]$. It is convenient to use the latter expression for values of C that do not correspond to integral values of T. This amounts to an interpolation. Thus, for general k, the expected working-set size is chosen to be C/k blocks (i.e., C pages), and we compute $\overline{M}^*[(\overline{S})^{*-1}(C/k)]$.

In Table 1 we give $\hat{T} = (\overline{S}^*)^{-1}(C/k)$ and $\overline{M}^*(\hat{T})$ for a fixed value of C and several blocking factors using the partitioning method described above, $\{\lambda_i\}$ of Fig. 2 and $r = 0.79$ as in Figs. 3 and 4. The results agree with observations that miss ratios for AAS are not greatly reduced by blocking pages together. (The lack of sensitivity to blocking shown in Table 1 is an inherent feature of the model. It can be shown [22] that for the partitioning method described above,

$$|\overline{M}[\overline{S}^{-1}(C)] - \overline{M}^*[(\overline{S}^*)^{-1}(C/k)]| \leq (1-r)\left(\frac{2k}{C} + \frac{33k^2}{C^2}\right).$$

Thus, if C/k, the expected working-set size (in *blocks*) is large, the miss ratios in the two cases are nearly equal.

Table 1 Effects of page blocking.

Expected working-set size C		*Number of pages per block, k* 1	4	8
5×10^4	$\hat{T}$	4.08×10^5	1.02×10^5	5.09×10^4
	$M(\hat{T})$	0.0989	0.0989	0.0989
2×10^5	$\hat{T}$	2.70×10^6	6.75×10^5	3.37×10^5
	$M(\hat{T})$	0.0479	0.0479	0.0479

Conclusions

A Markov chain model with a particularly simple structure has been studied and shown to provide analytic results for properties of interest in a reference string. A reasonable choice of parameters for the model yields results that agree well with observations of an AAS data string for large window sizes. However, reference strings from other data base systems should be studied to see if this model has general applicability. If so, then a model only slightly less tractable than the independent reference model shows promise of providing a useful tool in the study of data base systems.

Acknowledgments

The work described has benefited from numerous suggestions by A. C. McKellar. B. T. Bennett provided the LRU miss ratio values and other useful information concerning AAS. Discussions with R. Fagin and M. Z. Ghanem were helpful in clarifying several points. The presentation has benefited from a number of comments by the referees.

References

1. C. K. Chow, "On Optimization of Storage Hierarchies," *IBM J. Res. Develop.* **18,** 194 (May 1974).
2. J. Gecsei and J. A. Lukes, "A Model for the Evaluation of Storage Hierarchies," *IBM Syst. J.* **13,** 163 (1974).
3. E. Gelenbe, P. Tiberio, and J. C. A. Boekhorst, "Page Size in Demand-Paging Systems," *Acta Informat* **3,** 1 (1973).
4. D. E. Kuck and D. H. Lawrie, "The Use and Performance of Memory Hierarchies," *Computer and Information Sciences II*, edited by J. T. Lou, Academic Press, New York, 1969.
5. K. R. Kaplan and R. O. Winder, "Cache-Based Computer Systems," *Computer* **6,** 30 (March 1973).
6. J. H. Wimbrow, "A Large-Scale Interactive Administrative System," *IBM Syst. J.* **10,** 261 (1971).
7. A. V. Aho, P. J. Denning, and J. D. Ullman, "Principles of Optimal Page Replacement," *J. Assoc. Comput. Mach.* **18,** 80 (January 1971).
8. A. Chang, "Some Probabilistic Models of Storage Systems," *Research Report RJ 781*, IBM Research Division, San Jose, California, 1970.
9. P. J. Denning and S. C. Schwartz, "Properties of the Working-Set Model," *Comm. ACM* **15,** 191 (March 1972).
10. M. C. Easton and B. T. Bennett, "On Transient-Free Working-Set Statistics," *Research Report RC5140*, IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1970.
11. A. M. Yaglom, *Stationary Random Functions*, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1962.
12. E. G. Coffman, Jr., and P. J. Denning, *Operating System Theory*, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1973.
13. J. R. Spirn and P. J. Denning, "Experiments with Program Locality," *AFIPS Conf. Proc., Fall Jt. Comput. Conf.*, 1972, p. 611.
14. M. C. Easton, "Model for Interactive Data Base Reference String," *Research Report RC5050* (preliminary version), IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1974.
15. W. F. Frieberger, U. Grenander, and P. D. Sampson, "Patterns in Program References," *IBM J. Res. Develop.* **19,** 230 (1975).
16. J. W. Cohen, *The Single Server Queue*, John Wiley and Sons, Inc., New York, 1969.
17. M. Z. Ghanem and H. Kobayashi, "A Parametric Representation of Program Behavior in a Virtual Memory System," *Research Report RC4560*, IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1973.

18. D. E. Knuth, *The Art of Computer Programming*, Vol. 3, Addison-Wesley Publishing Company, Reading, Massachusetts, 1973.
19. B. T. Bennett, IBM Thomas J. Watson Research Center, Yorktown Heights, New York, private communication.
20. P. C. Yue and C. K. Wong, "On the Optimality of the Probability Ranking Scheme in Storage Applications," *J. Assoc. Comput. Mach.* **20,** 624 (October 1973).
21. C. H. Hardy, J. E. Littlewood, and G. Polya, *Inequalities*, Cambridge University Press, Cambridge, England, 1964.
22. R. Fagin and M. C. Easton, "The Independence of Miss Ratio on Page Size," *Research Report RC5006* (preliminary version), IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1974 (to appear in *J. Assoc. Comput. Mach.*).

Received April 28, 1975; revised July 25, 1975

The author is located at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York 10598.