A Counterintuitive Example of Computer Paging

Ronald Fagin
IBM Thomas J. Watson Research Center

A counterexample is exhibited to a natural conjecture concerning the optimal way to group records into pages in the independent reference model of computer paging (an organization is said to be optimal if the "least recently used" miss ratio is minimized).

Key Words and Phrases: least recently used, most likely together, independent reference model, storage organization, record allocation

CR Categories: 3.70, 3.73, 4.0, 5.30

1. Introduction

Consider the following problem: given a set of equal-sized records and their request probabilities, how should we partition the records into equal-sized pages so as to minimize the expected LRU ("least recently used") miss ratio [1,6] in a paged computing system with two-level storage hierarchy? We will show that in a very simple model of record references (the "independent reference model" of [1,5]), a natural conjecture is false.

Assume that there are mk records $R_1, \ldots, R_{mk}$, and that at each discrete time unit exactly one record is referenced; record i is referenced at time t with probability $p_i (1 \leq i \leq mk)$, independent of past history. Of course $\sum p_i = 1$. We wish to distribute the mk records among m pages with exactly k records to a page. Under the LRU memory management policy, if there is a page fault, that is, if a page is referenced that is not in main memory, then that page is moved into main memory and the page that has been least recently referenced is removed. The *expected LRU miss ratio M* is the limiting probability of a page fault. The expected LRU hit ratio (which is $1 - M$) has a natural interpretation in the independent reference model as the expected weight of main memory. We wish to know which partitioning of records into pages minimizes the expected LRU miss ratio.

For ease in notation, assume that $p_1 \geq p_2 \geq \cdots \geq p_{mk} \geq 0$. A natural conjecture is that the optimal organization is obtained by placing $R_1, \ldots, R_k$ on one page, $R_{k+1}, \ldots, R_{2k}$ on another page, and so on. In other words, the most frequently referenced records are assigned to one page, the next most frequently referenced records are assigned to another page, and so on. We will call this particular allocation of records to pages the "*most likely together*" organization. (Actually, this organization need not be unique if the p_i's are not all different. For example, if $p_4 = p_5$, then we could interchange the roles of records R_4 and R_5. However, this need not concern us.) Yue and Wong [7] show that the most likely together organization is best under two criteria of optimality: minimal expected working set size [3], and minimal expected LRU stack distance [6]. But we will show that this organization does not always minimize the expected LRU miss ratio.

2. A Class of Counterexamples

Assume that there are eight records $R_1, \ldots, R_8$; that we wish to allocate two records to a page; and that first-level memory is just large enough to hold six records (three pages). Assume that records R_1 and R_2 each have access probability p; that records R_3, R_4, R_5, and R_6 each have access probability r; that record R_7 has access probability $r - \epsilon$; and that record R_8 has access probability ϵ. Thus, $2p + 5r = 1$. Assume that

Author's present address: IBM Research Laboratory, 5600 Cottle Road-K53/282, San Jose, CA 95193.

$p > r > \epsilon \geq 0$. For simplicity, the reader can assume throughout that $\epsilon = 0$. We are allowing the possibility that $\epsilon \neq 0$ so that no record need have access probability zero.

In the most likely together organization, records R_1 and R_2 will be on one page, R_3 and R_4 on another page, and so on. So there will be four pages, with page reference probabilities $2p$, $2r$, $2r$, and r. If we apply King's formula [5; 2, p. 272] for the expected LRU miss ratio, we find, after considerable computation[1] that the miss ratio is

$$\frac{2p(1-2p)(1237+5479p-6812p^2+4416p^3+9792p^4)}{75(2+p)(3+4p)(1+3p)(1+8p)}. \tag{1}$$

We now consider another method of allocating records to pages. Page 1 contains records R_1 and R_4 ; page 2 contains R_2 and R_3 ; page 3 contains R_5 and R_6 ; and page 4 contains R_7 and R_8 . So the page reference probabilities are $(p + r)$, $(p + r)$, $2r$, and r. Again applying King's formula, we find that the miss ratio is

$$\frac{(1-2p)(1+3p)(1188+1131p-3973p^2-1196p^3+228p^4)}{75(2+p)(3+4p)(3-p)(4-3p)}. \tag{2}$$

Expression (1) minus expression (2) turns out to be

$$\frac{(1-2p)(1-7p)^2(-1188-4707p+1207p^2+7932p^3-2916p^4+864p^5)}{75(2+p)(3+4p)(1+3p)(1+8p)(3-p)(4-3p)}. \tag{3}$$

We are interested in whether expression (3) is positive or negative. Clearly, (3) is positive, for $0 < p < \frac{1}{2}$, iff the fifth-degree polynomial which is a subexpression of (3) is positive, namely:

$$-1188-4707p+12071p^2+7932p^3-2916p^4+864p^5. \tag{4}$$

It is straightforward to check that (4) has a (nonmultiple) zero at $\zeta = .4729\ldots$, and that (4) is positive for $p > \zeta$. Hence, if $\zeta < p < \frac{1}{2}$, then miss ratio (1) is larger than miss ratio (2). So, in this case, the most likely together organization is not optimal.

As an example, let $p = .49$ and $\epsilon = .001$. Then there are eight records, with request probabilities .49, .49, .004, .004, .004, .004, .003, and .001. Assume as before that main memory is just large enough to hold six records, or three pages of two records each. When we allocate the records to pages, with two records to a page, the miss ratio under the most likely together organization is .005868 .... Under the unusual organization we described, where records R_1 and R_4 are in one page, R_2 and R_3 in another page, R_5 and R_6 in another page, and R_7 and R_8 in the final page, the miss ratio is .005548 ..., which is smaller.

[1] All symbolic computations (which would have been excessively tedious to carry out by hand) were executed by the Scratchpad system [4] on the IBM 370/158 at the IBM Thomas J. Watson Research Center.

3. Summary

We have shown that in the independent reference model, the natural method of grouping records into pages, in which the most frequently referenced records are allocated to the same page, does not necessarily minimize the expected LRU miss ratio. Whether this organization is near-optimal is an open problem.

Received October 1974; revised March 1975

References

1. Aho, A.V., Denning, P.J., and Ullman, J.D. Principles of optimal page replacement. *J. ACM 18*, 1 (Jan. 1971), 80–93.
2. Coffman, E.G., Jr., and Denning, P.J. *Operating Systems Theory*. Prentice-Hall, Englewood Cliffs, N.J., 1973.
3. Denning, P.J. The working set model for program behavior. *Comm. ACM 11*, 5 (May 1968), 323–333.
4. Griesmer, J.H., and Jenks, R.D., Experience with an on-line symbolic mathematics system. Proceedings of the ONLINE72 Conference, Brunel University, Uxbridge, Middlesex, England, September 4–7, 1972. (Also available as The SCRATCHPAD system, IBM Research Rep. RC 3925, Yorktown Heights, N.Y.
5. King, W.F. III. Analysis of demand paging algorithms. Proc. IFIP Congress 71, Ljubljana, Yugoslavia. North-Holland Pub. Co., Amsterdam, 1971, pp. 485–490.
6. Mattson, R., Gecsei, J., Slutz, D., and Traiger, I. Evaluation techniques for storage hierarchies. *IBM Syst. J. 9*, 2 (1970), 78–117.
7. Yue, P.C. and Wong, C.K. On the optimality of the probability ranking scheme in storage applications. *J. ACM 20*, 4 (Oct. 1973), 624–633.