

LRU IS BETTER THAN FIFO UNDER THE INDEPENDENT REFERENCE MODEL

J. VAN DEN BERG,* *CWI*
A. GANDOLFI,** *Courant Institute*

Abstract

Consider a two-level storage system operating with the least recently used (LRU) or the first-in, first-out (FIFO) replacement strategy. Accesses to the main storage are described by the independent reference model (IRM). Using the FKG inequality, we prove that the miss ratio for LRU is smaller than or equal to the miss ratio for FIFO.

REPLACEMENT ALGORITHMS; FKG INEQUALITY; MEMORY MANAGEMENT; MISS RATIO; FIFO; LRU

AMS 1991 SUBJECT CLASSIFICATION: PRIMARY 60C05
SECONDARY 90B05; 60K99

1. Introduction

Consider a storage system which consists of two levels, named main storage and secondary storage. Suppose there are n items, denoted by $1, 2, \dots, n$, which can be located in secondary or main storage. The main storage can contain at most m items. Items never leave the system and no new items enter the system. If an item is required, first the main storage is inspected. If it is not present there, it is taken from secondary storage and put in main storage. If the main storage is already full, an item is removed to secondary storage according to some replacement algorithm. This situation occurs for instance in computer systems with paged main and secondary memory (see e.g. Matick (1977)).

We restrict our attention to the well-known least recently used (LRU) and first-in, first-out (FIFO) algorithms. In the case of FIFO, that item is replaced which has been in main storage for the longest time. In the case of LRU that item is replaced whose last reference was earliest among those items in main storage. Furthermore we assume that

Received 5 November 1990.

* Postal address: CWI, Kruislaan 413, 1098 SJ, Amsterdam, The Netherlands.

Research carried out while the author was visiting Cornell University, Ithaca, NY, and partly supported by the US Army Research Office.

** Present address: Department of Statistics, 367 Evans Hall, University of California, Berkeley, CA 94720, USA.

Research carried out while the author was visiting the University of California at Davis.

the string of page references is a sequence of i.i.d. random variables. This is known as the independent reference model (IRM). There is much literature on these systems; for references see Aven et al. (1987), Chapters 4 and 5.

We are interested in the miss ratio, which is defined as the steady-state probability that the item currently required is not in main storage (in computer storage terms, this is the frequency of page faults). It is widely believed that, under the IRM, the miss ratio for LRU is smaller than or equal to the miss ratio for FIFO (see Aven et al. (1987), who formulate this as a special case of a more general monotonicity conjecture for an indexed family of replacement algorithms; another special case of that conjecture has been proved by Kan and Ross (1980)). We give a proof based on the FKG inequality.

2. Formal statement of the result

Consider the IRM model described in Section 1. Let $S = \{1, \dots, n\}$ be the set of items, $\mathbf{p} = (p_1, p_2, \dots, p_n)$ their reference probabilities and m the size of main storage. To avoid trivialities we assume that p_1, \dots, p_n are strictly positive and that $m \leq n$. The miss ratio, introduced before, depends on \mathbf{p} , m and the replacement algorithm A , and is denoted by $F_m(A, \mathbf{p})$. When the context is clear, we omit one or more of the parameters.

Let $\{\xi_k\}_{k=0}^\infty$ be a sequence of independent and identically distributed random variables with $P(\xi_0 = j) = p_j, j = 1, 2, \dots, n$. When we regard ξ_0 as the item currently being referenced and ξ_k the k th item previously referenced, $k = 1, \dots$, it is not difficult to see that for LRU the miss ratio, defined in the introduction, is equal to

$$(1) \quad F_m(\text{LRU}) = P(\xi_0 \text{ is not equal to one of the first } m \text{ distinct values in the sequence } \xi_1, \xi_2, \dots).$$

Further, it is clear that, for distinct $x_1, \dots, x_m \in S$,

$$(2) \quad \begin{aligned} &P(\text{the first } m \text{ distinct values in the sequence } \xi_1, \xi_2, \dots \text{ are, in order of} \\ &\text{occurrence, equal to } x_1, x_2, \dots, x_m) \\ &= p_{x_1} \frac{p_{x_2}}{1 - p_{x_1}} \frac{p_{x_3}}{1 - p_{x_1} - p_{x_2}} \dots \frac{p_{x_m}}{1 - p_{x_1} - \dots - p_{x_{m-1}}}. \end{aligned}$$

Combining (1) and (2) yields the following result by King (1971):

$$(3) \quad F_m(\text{LRU}) = \sum_{(x_1, \dots, x_m) \in \Lambda} \frac{p_{x_1} p_{x_2} \dots p_{x_m} (1 - p_{x_1} - p_{x_2} - \dots - p_{x_m})}{(1 - p_{x_1})(1 - p_{x_1} - p_{x_2}) \dots (1 - p_{x_1} - p_{x_2} - \dots - p_{x_{m-1}})},$$

where the sum is over the set

$$(4) \quad \Lambda = \{(x_1, \dots, x_m) \in S^m : x_1, \dots, x_m \text{ are distinct}\}.$$

As to FIFO, it has been shown by King (1971) and Aven and Sokolov (1971) that

$$(5) \quad F_m(\text{FIFO}) = \frac{\sum_{(x_1, \dots, x_m) \in \Lambda} p_{x_1} p_{x_2} \dots p_{x_m} (1 - p_{x_1} - p_{x_2} - \dots - p_{x_m})}{\sum_{(x_1, \dots, x_m) \in \Lambda} p_{x_1} p_{x_2} \dots p_{x_m}}.$$

See also Aven et al. (1987), Section 4.6, for a discussion of the above results. We prove the following theorem.

Theorem. Under the independent reference model $F_m(\text{LRU}) \leq F_m(\text{FIFO})$.

The proof is based on the FKG inequality (Fortuin et al. (1971)). This inequality is as follows.

Let Ω be a finite set, $<$ a partial order on Ω and μ a weight function (measure) on Ω . We say that $(\Omega, <, \mu)$ satisfies the FKG conditions if, for all $x, y, z \in \Omega$:

(i) there exist $x \wedge y$ and $x \vee y$, which are (respectively) the largest (smallest) element of Ω smaller (larger) than x and y ;

(ii) $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$;

(iii) $\mu(x) \geq 0$; $\mu(x)\mu(y) \leq \mu(x \wedge y)\mu(x \vee y)$.

A function f on Ω is *increasing* if $x < y$ implies $f(x) \leq f(y)$.

The FKG inequality says that if $(\Omega, <, \mu)$ satisfies the FKG conditions, then, for all increasing functions f and g ,

$$(6) \quad \sum_{x \in \Omega} \mu(x) f(x) g(x) \sum_{y \in \Omega} \mu(y) \geq \sum_{x \in \Omega} \mu(x) f(x) \sum_{y \in \Omega} \mu(y) g(y).$$

It should be noted that (i)–(iii) are sufficient but not necessary conditions for (6) to hold. The conditions (i) and (ii) together are equivalent to saying that $(\Omega, <)$ is a distributive lattice. The FKG inequality has been extensively used in statistical mechanics and related fields. A different type of application is L. Shepp’s proof of the XYZ conjecture (Shepp (1982)). A well-known special case is that where Ω is a product of totally ordered finite sets, μ is a product measure on Ω and $<$ is the natural partial order on Ω . This case is also known as Harris’s inequality (Harris (1960)).

3. Proof of the theorem

Without loss of generality we may assume that $p_1 \leq p_2 \leq \dots \leq p_n$.

For $x = (x_1, \dots, x_m) \in S^m$ let $\mu(x) = p_{x_1} p_{x_2} \dots p_{x_m}$, and define $x < y$ iff $x_i \leq y_i$, $i = 1, \dots, m$.

It is easy to verify that $(S^m, <, \mu)$ satisfies the FKG conditions. In particular we have $x \wedge y = (\min(x_1, y_1), \dots, \min(x_m, y_m))$ and $x \vee y = (\max(x_1, y_1), \dots, \max(x_m, y_m))$. In fact this is the special case corresponding to Harris’s inequality, mentioned at the end of Section 2. However, this special case is not suitable for our purpose. We are essentially interested in Λ (see (4)), but Λ , with the partial order defined above, clearly violates (i) in the FKG conditions. In spite of this fact, the following lemma concerning Λ holds.

Lemma. Let f and g be functions on Λ and let f or g be permutation invariant (i.e. independent of the order of its arguments). If f and g are increasing, then

$$(7) \quad \sum_{x \in \Lambda} \mu(x) f(x) g(x) \sum_{y \in \Lambda} \mu(y) \geq \sum_{x \in \Lambda} \mu(x) f(x) \sum_{y \in \Lambda} \mu(y) g(y).$$

Equivalently, if f is decreasing and g increasing,

$$\sum_{x \in \Lambda} \mu(x) f(x) g(x) \sum_{y \in \Lambda} \mu(y) \leq \sum_{x \in \Lambda} \mu(x) f(x) \sum_{y \in \Lambda} \mu(y) g(y).$$

Proof of the lemma. We prove the first part of the lemma. The second part then follows directly after replacing f by $-f$. So let f and g be increasing and suppose at least one of them, say f , is permutation invariant. Let Π be the set of all permutations of $(1, \dots, m)$. For $\pi \in \Pi$ and $x = (x_1, \dots, x_m) \in S^m$ define $\pi(x) = (x_{\pi(1)}, x_{\pi(2)}, \dots, x_{\pi(m)})$.

Let the set Ω be defined by

$$\Omega = \{(x_1, \dots, x_m) \in S^m : x_1 < x_2 < \dots < x_m\}.$$

Since $\Omega \subset S^m$ we may consider the restriction of μ and $<$ to Ω . Since we already know that $(S^m, <, \mu)$ satisfies the FKG conditions, the only thing we have to do to ensure that $(\Omega, <, \mu)$ satisfies the FKG conditions is to check that Ω is closed under \wedge and \vee . So assume $x, y \in \Omega$ and $1 \leq i < j \leq m$. Then $(x \vee y)_j \geq x_j > x_i$ and $(x \vee y)_j \geq y_j > y_i$, hence $(x \vee y)_j > \max(x_i, y_i) = (x \vee y)_i$. Similarly, $(x \wedge y)_i \leq x_i < x_j$ and $(x \wedge y)_i \leq y_i < y_j$, hence $(x \wedge y)_i < \min(x_j, y_j) = (x \wedge y)_j$. We conclude that Ω is indeed closed under \wedge and \vee and $(\Omega, <, \mu)$ satisfies the FKG conditions.

The inequality (7) can now easily be derived:

$$\begin{aligned} \sum_{x \in \Lambda} \mu(x) f(x) g(x) \sum_{y \in \Lambda} \mu(y) &= \sum_{x \in \Omega} \sum_{\pi \in \Pi} \mu(\pi(x)) f(\pi(x)) g(\pi(x)) \sum_{y \in \Omega} \sum_{\pi \in \Pi} \mu(\pi(y)) \\ &= \sum_{x \in \Omega} \mu(x) f(x) \sum_{\pi \in \Pi} g(\pi(x)) m! \sum_{y \in \Omega} \mu(y) \\ &\geq m! \sum_{x \in \Omega} \mu(x) f(x) \sum_{y \in \Omega} \mu(y) \sum_{\pi \in \Pi} g(\pi(y)) \\ &= \sum_{x \in \Lambda} \mu(x) f(x) \sum_{y \in \Lambda} \mu(y) g(y). \end{aligned}$$

The first equality follows from the definition of Ω and Λ , the second from the permutation invariance of f and μ . The inequality follows from the FKG inequality (6) (note that, since g is increasing, for every π the function $y \rightarrow g(\pi(y))$ is increasing and hence the function $y \rightarrow \sum_{\pi \in \Pi} g(\pi(y))$ is increasing). The last equality follows again from the definition of Λ and Ω and the permutation invariance of f and μ .

Proof of the theorem. The theorem follows now as an application of the lemma: note that (3) can be written as follows:

$$(8) \quad F_m(\text{LRU}) = \sum_{x \in \Lambda} \mu(x) f(x) g(x),$$

where the functions f and g are defined by

$$f(x_1, \dots, x_m) = 1 - p_{x_1} - p_{x_2} - \dots - p_{x_m},$$

and

$$g(x_1, \dots, x_m) = \frac{1}{(1 - p_{x_1})(1 - p_{x_1} - p_{x_2}) \cdots (1 - p_{x_1} - \dots - p_{x_{m-1}})}.$$

It is immediately clear that f is decreasing and permutation invariant and that g is increasing.

Applying the above lemma to (8) yields:

$$F_m(\text{LRU}) \sum_{y \in \Lambda} \mu(y) \leq \sum_{x \in \Lambda} \mu(x) f(x) \sum_{y \in \Lambda} \mu(y) g(y).$$

Note that, by (2) and the definition of μ and g ,

$$\sum_{y \in \Lambda} \mu(y) g(y) = P(\text{the sequence } \xi_1, \xi_2, \dots \text{ contains at least } m \text{ different items}) = 1.$$

Hence,

$$F_m(\text{LRU}) \leq \frac{\sum_{x \in \Lambda} \mu(x) f(x)}{\sum_{x \in \Lambda} \mu(x)},$$

which, by (5) and the definition of μ and f , equals $F_m(\text{FIFO})$.

References

- AVEN, O. I. AND SOKOLOV, V. B. (1971) On some methods of memory management (in Russian) *Tech. Cybernet.* **2**, 89–94.
- AVEN, O. I., COFFMAN, E. G., JR., AND KOGAN, Y. A. (1987) *Stochastic Analysis of Computer Storage*. Reidel, Dordrecht.
- VAN DEN BERG, J. AND TOWSLEY, D. (1990) Properties of the miss ratio for a 2-level storage model with LRU or FIFO replacement strategy and independent references. Preprint.
- FORTUIN, C. M., KASTELEYN, P. W. AND GINIBRE, J. (1971) Correlation inequalities on some partially ordered sets. *Comm. Math. Phys.* **22**, 89–103.
- HARRIS, T. E. (1960) A lower bound for the critical probability in a certain percolation process. *Proc. Camb. Phil. Soc.* **56**, 13–20.
- KING, W. F. (1971) Analysis of paging algorithms. *Proc. IFIP Congress*, Lyublyana, Yugoslavia, August 1971, 485–490.
- KAN, Y. C. AND ROSS, S. M. (1980) Optimal list ordering under partial memory constraints. *J. Appl. Prob.* **17**, 1004–1015.
- MATICK, R. E. (1977) *Computer Storage and Technology*. Wiley, New York.
- SHEPP, L. (1982) The XYZ conjecture and the FKG inequality. *Ann. Prob.* **10**, 824–827.