HEIRARCHICAL REPLACEMENT DECISIONS
IN HIERARCHICAL STORES⁺

Ozalp Babaoglu*
Domenico Ferrari**

TR 82-512
August 1982

Department of Computer Science
Upson Hall
Cornell University
Ithaca, New York 14853

⁺Partial support for this work was provided by the National Science Foundation under Grant No. MCS78-24618 and the Italian National Research Council through a Fellowship for Foreign Scholars to the first author.

*Department of Computer Science, Cornell University, Ithaca, NY 14853

**Computer Science Division, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720

Hierarchical Replacement Decisions in Hierarchical Stores †

Özalp Babaoğlu ‡

Domenico Ferrari *

Abstract

One of the primary motivations for implementing virtual memory is its ability to automatically manage a hierarchy of storage systems with different characteristics. The composite system behaves as if it were a single-level system having the more desirable characteristics of each of its constituent levels. In this paper we extend the virtual memory concept to within each of the levels of the hierarchy. Each level is thought of as containing two additional levels within it. This hierarchy is not a physical one, but rather an artificial one arising from the employment of two different replacement algorithms. Given two replacement algorithms, one of which has good performance but high implementation cost and the other poor performance but low implementation cost, we propose and analyze schemes that result in an overall algorithm having the performance characteristics of the former and the cost characteristics of the latter. We discuss the suitability of such schemes in the management of storage hierarchies that lack page reference bits.

Index Terms

Virtual memory, replacement algorithms, FIFO, LRU, working set, IRM, Markov chain, performance evaluation.

† Partial support for this work was provided by the National Science Foundation under Grant No. MCS78-24618 and the Italian National Research Council through a Fellowship for Foreign Scholars to the first author.

‡ Department of Computer Science, Cornell University, Ithaca, New York 14853.

* Computer Science Division, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, California 94720.

1. Introduction

Economic realities force the construction of large virtual address spaces in the form of hierarchies. A typical two-level hierarchy consists of a fast, expensive, and small (the latter two being a consequence of the first attribute) primary memory in addition to a slow, cheap and large secondary memory. Obtaining the appearance of a single level that is fast and large by the automatic management of the hierarchy is probably one of the primary motivations for incorporating virtual memory in a computer system [15, 17].

Every level except the last (i.e., the slowest) in a hierarchy adopts a replacement policy to remove data down to the next level when necessary. The topic of replacement algorithms that are suitable for such environments has been the subject of a large number of investigations [23]. Specifically, the virtual memory system we are concerned with employs *demand paging* [11]. The replacement algorithms that have been proposed for this environment can be crudely classified as being either inexpensive to implement but having poor performance, or being expensive to implement but having better performance. Our informal notion of cost of implementation considers factors such as hardware support other than the basic address translation mechanism and the complexity of the associated software. Our criterion for performance, on the other hand, is based on the frequency of page faults. Examples of replacement algorithms of the cheap-to-implement-but-poor-performance type include the First-In-First-Out (FIFO) and the Random (RAND) algorithms [6]. The Least Recently Used (LRU) [20] and the Working Set (WS) [10] algorithms, on the other hand, are examples of the expensive-to-implement-but-good-performance class.

In the next few sections, we introduce a class of *hybrid* replacement policies that combine two algorithms, one of each of the above categories, within a single level of the physical storage hierarchy. After deriving expressions for their performance, we demonstrate that these algorithms in fact achieve performances close to those of LRU and WS while having implementation costs comparable to those of FIFO and RAND. The next section introduces the program model on which our analyses will be based.

2. The Independent Reference Model

The mathematical analysis of a replacement algorithm requires a model of the programs on which the policy operates. For our purposes, an execution of a program consisting of n pages labeled $\{1, 2, \dots, n\}$ results in a page reference string, $r_1, r_2, r_3, \dots, r_{t-1}, r_t, r_{t+1}, \dots$, where $r_t = i$ if page i is referenced at time instant t (memory references are assumed to occur at equidistant time points, and we define their distance to be the unit of time). We will assume a particularly simple stochastic structure for the reference string, known as the Independent Reference Model (IRM) [1]. As the name implies, the string $\{r_i; i=1, 2, \dots\}$ is assumed to be a sequence of independent, identically distributed random variables from the population $\{1, 2, \dots, n\}$ where $\Pr(r_t = i) = \beta_i$ for all t and $\sum_{i=1}^n \beta_i = 1$.

Modeling an actual program with the IRM involves obtaining point estimates for the model parameters (β_i 's) that are simply the frequency counts of pages in an actual reference string generated by the program. Baskett and Rafii have proposed an alternate method for obtaining estimates for the β_i 's, called the A_0 *Inversion Model* [5]. The resulting model, although structurally identical to the IRM, has much better predictive capabilities for the performance of real programs.

3. Hybrid Policies

In a demand-paged virtual memory system, referencing a page that is *invalid* — not in main memory — causes a trap, which is known as a *page fault*. We note that, even in the absence of any additional hardware support, this address translation mechanism can be put to use to detect references to pages that are already in memory. All that is required is that we be able to distinguish these faults from normal page faults and refrain from initiating the I/O operation. This special state of a page will be called the *reclaimable* state, and will be identified by one additional bit in each page table entry. Since this method of detecting references to pages comes at a cost (to be discussed later), we are interested in replacement algorithms that collect reference informa-

tion only for a subset of the pages that a program has in memory. More formally, we have partitioned the set of pages in memory into two disjoint classes, $\{valid\}$ and $\{reclaimable\}$, such that

$$\{memory\} = \{valid\} \cup \{reclaimable\} \text{ and } \{valid\} \cap \{reclaimable\} = \emptyset.$$

To keep the cost of generating spurious faults to the reclaimable pages at reasonably low levels, we would like to have

$$|\{reclaimable\}| \ll |\{valid\}|.$$

We make these statements more precise in the following sections. For reasons which will become clear below, we shall refer to the set $\{valid\}$ as **T** (for “top”), and to the set $\{reclaimable\}$ as **B** (for “bottom”). Having partitioned the program’s pages in memory into these two classes, we consider various policies for top-to-bottom and memory-to-secondary storage replacements. Since in our investigation we assume a single program to be executing, the analyses presented in the forthcoming sections consider *local* management policies. Extensions of these policies to *global* replacement schemes in multiprogramming environments are discussed in section 6.

The two reasonable choices for the management of **T** are FIFO and RAND, since for those pages we have no reference information. Algorithms such as the Sampled Working Set [14, 21] and Clock [8, 12] policies that require reference information, are clearly out of consideration in this region. On the other hand, we can employ either the LRU or the WS algorithm for the management of **B** since the necessary information can be gathered at every occurrence of an artificial page fault. This results in four possible combinations that make up the hybrid class to be studied. Denoted $H_{FIFO-LRU}$, $H_{RAND-LRU}$, $H_{FIFO-WS}$ and $H_{RAND-WS}$, they are the FIFO-LRU, RAND-LRU, FIFO-WS and the RAND-WS hybrid policies, respectively.

3.1. Fixed Partition Hybrids

Employment of the LRU policy for the management of **B** results in hybrid algorithms that operate in a fixed size memory partition. However, in a multiprogramming environment, the use of a common bottom amongst all the active programs results in a variable size partition for each

even though the tops are strictly local. We comment about such extensions in section 6 and restrict our study here to uniprogramming environments. For the following analysis, assume that $\mathbf{T}$ consists of k pages (i.e., $|\mathbf{T}| = k$), where k is the parameter of the policy, whereas the fixed partition size is m pages (i.e., $|\{\text{memory}\}| = m$, $k \leq m \leq n$).

3.1.1. The FIFO-LRU Hybrid Policy

Given a page reference, r_t , at time t , the operation of the $H_{FIFO-LRU}(k)$ policy is as follows:

- H1: If $r_t \in \mathbf{T}$, no control state change takes place. This is because this type of reference is transparent to our mechanism.
- H2: (Reclaim) If $r_t \in \mathbf{B}$, then $\mathbf{T} \leftarrow \mathbf{T} + r_t - i$, where i is the FIFO (oldest) page in $\mathbf{T}$, and $\mathbf{B} \leftarrow \mathbf{B} - r_t + i$. Note that, in these expressions, “+” and “-” denote set membership operations.
- H3: (Page fault) If $r_t \notin \{\text{memory}\}$, then $\mathbf{T} \leftarrow \mathbf{T} + r_t - i$, where page i is as in H2, and $\mathbf{B} \leftarrow \mathbf{B} + i - j$ where page j is the one that has been (approximately) least recently used amongst all pages in $\mathbf{B}$.

We cannot state that page j is exactly the LRU page because the ordering amongst the top is by time of entry and not by recency of use. Consequently, there may be pages in memory that have been referenced earlier than page j if, for example, page j was referenced just prior to its most recent departure from $\mathbf{T}$. A more appropriate name for the replacement policy employed in $\mathbf{B}$ is *Least Recently Reclaimed*. In section 4 we present numerical results which suggest that, under a wide range of circumstances, the page replaced by these fixed partition hybrid policies from the bottom is very close to being the LRU page.

If we envision the control state associated with the algorithm as constituting a stack, the $H_{FIFO-LRU}$ policy can be regarded as a modification to the pure LRU policy where references to the top k positions of the LRU stack cause no control state change [20]. Note that, for the degenerate case $k=1$, the page replaced from memory to secondary storage by the $H_{FIFO-LRU}$

policy is exactly the same page that would be replaced by the pure LRU policy. Furthermore, when $k=m$, the $H_{FIFO-LRU}$ policy degenerates into the pure FIFO policy. We point out that, as presented above, the $H_{FIFO-LRU}$ algorithm is identical to the *Segmented FIFO* (SFIFO) algorithm as described by Turner and Levy [24], and to the A_k^1 algorithm analyzed by Aven *et al* [2].

The performance index that we will use to compare different policies is the *steady-state fault rate*. For a replacement algorithm A , the steady-state fault rate is defined as:

$$F(A) = \lim_{t \rightarrow \infty} [\mathbf{Pr}(r_t \notin \{\text{memory}\})].$$

In other words, $F(A)$ is the limiting probability with which a reference to a page causes a page fault.

We are now in a position to derive an expression for $F(H_{FIFO-LRU}(k))$, based on the IRM. Such an expression has been previously obtained by Aven *et al* [2]. Here we show our derivation to illustrate the analysis technique which is similar to that used in [7] and which will be applied to the other hybrid algorithms as well. Let $\mathbf{s}=[j_1, j_2, \dots, j_k, j_{k+1}, \dots, j_m]$ be an m -tuple (without repetitions) corresponding to the memory control state of the policy. The first k entries of $\mathbf{s}$ contain the page names that constitute $\mathbf{T}$, whereas the remaining $m-k$ entries contain the names of the elements of $\mathbf{B}$. Define the Markov chain $\{X_t, t=0, 1, \dots\}$, with $X_t=\mathbf{s}$ if the memory control state at time t is given by $\mathbf{s}$. Let $\mathbf{Q}=\{\mathbf{s}\}$ denote the state space of this Markov chain. From the description of $\mathbf{s}$, one may conclude that $\mathbf{Q}$ consists of the set of all permutations of m elements chosen from n items. Therefore,

$$|\mathbf{Q}| = \binom{n}{m} m! = n! / (n-m)!.$$

The one-step transition probabilities denoted by

$$p(\mathbf{s}, \mathbf{s}') = \mathbf{Pr}(X_t = \mathbf{s}' \mid X_{t-1} = \mathbf{s}), \quad t \geq 1$$

can be determined easily based on the IRM parameters and on the algorithm's description. Specifically, for $H_{FIFO-LRU}$,

$$p(\mathbf{s}, \mathbf{s}') = \begin{cases} \sum_{i=1}^k \beta_{j_i}, & \text{if } \mathbf{s}' = \mathbf{s} \\ \beta_{j_l}, & \text{if } \mathbf{s}' = [j_l, j_1, j_2, \dots, j_{l-1}, j_{l+1}, \dots, j_m], \quad k < l \leq m \\ \beta_j, & \text{if } \mathbf{s}' = [j, j_1, j_2, \dots, j_{m-2}, j_{m-1}], \quad j \notin \mathbf{s} \\ 0, & \text{otherwise.} \end{cases}$$

The three nonzero cases correspond to the H1, H2, and H3 events of the algorithm's description, respectively. As the above-defined chain is ergodic, the limiting state occupancy probabilities, π , exist and satisfy

$$\pi = \pi \cdot \mathbf{P}, \quad (1)$$

where $\mathbf{P} = [p(\mathbf{s}, \mathbf{s}')]$ is the one-step transition matrix [22]. The limiting state occupancy probabilities, π , which are the eigenvalues of $\mathbf{P}$, have also to satisfy the normalization condition $\sum_{\mathbf{s} \in \mathcal{Q}} \pi_{\mathbf{s}} = 1$. For a particular state $\mathbf{s} = [j_1, j_2, j_3, \dots, j_m]$, the matrix equation (1) can be written as

$$\pi_{\mathbf{s}} = \pi_{\mathbf{s}} \sum_{i=1}^k \beta_{j_i} + \beta_{j_1} \left[\sum_{j \notin \mathbf{s}} \pi_{\mathbf{u}_j} + \sum_{k < l \leq m} \pi_{\mathbf{v}_l} \right] \quad (2)$$

where $\mathbf{u}_j = [j_2, j_3, \dots, j_m, j]$ and $\mathbf{v}_l = [j_2, j_3, \dots, j_l, j_1, j_{l+1}, \dots, j_m]$. Note that states $\mathbf{u}_j$ and $\mathbf{v}_l$ have been constructed so that a reference to page j_1 causes the algorithm to make a transition to state $\mathbf{s}$.

Lemma 1: For the $H_{FIFO-LRU}$ policy with parameter k , the equilibrium probability of state $\mathbf{s} = [j_1, j_2, j_3, \dots, j_m]$ is given by

$$\pi_{\mathbf{s}} = \frac{\prod_{i=1}^m \beta_{j_i}}{G(k) \prod_{i=2}^{m-k+1} D_i(\mathbf{s})}$$

where

$$G(k) = \frac{(n-m)!}{(n-k)!} \sum_{\mathbf{s} \in \mathbf{Q}} \prod_{i=1}^k \beta_{j_i} \quad \text{and} \quad D_i(\mathbf{s}) = 1 - \sum_{l=1}^{m-i+1} \beta_{j_l}.$$

The proof is presented in the Appendix.

That the normalizing condition $\sum_{\mathbf{s} \in \mathbf{Q}} \pi_{\mathbf{s}} = 1$ is satisfied can be shown by an aggregation argument where first the $\pi_{\mathbf{s}}$'s are summed over the $(n-k)!/(n-m)!$ states that have the same $\mathbf{T}$ (the first k elements of $\mathbf{s}$) and then the resulting aggregates are summed to cover the entire state space $\mathbf{Q}$. The second summation can be shown to be the normalization condition for the equilibrium probabilities of the memory control states for a k -page memory managed by the pure FIFO policy which we know to hold.

Having the above Lemma at hand, the following theorem can easily be proved:

Theorem 1: The steady-state fault rate generated by the $H_{FIFO-LRU}$ policy with parameter k and operating in a memory of m page frames is given by

$$F(H_{FIFO-LRU}(k)) = \sum_{\mathbf{s} \in \mathbf{Q}} \left[G^{-1}(k) D_1^2(\mathbf{s}) \frac{\prod_{i=1}^m \beta_{j_i}}{\prod_{i=1}^{m-k+1} D_i(\mathbf{s})} \right]$$

where $G(k)$ and $D_i(\mathbf{s})$ are as defined in Lemma 1.

Proof. Note that, given the current memory state $\mathbf{s}$, the probability of a page fault is simply $1 - \sum_{i=1}^m \beta_{j_i}$ which is by definition $D_1(\mathbf{s})$. Thus, conditioning on the state $\mathbf{s}$, we can write

$$\mathbf{Pr}(\text{page fault}) = \sum_{\mathbf{s} \in \mathbf{Q}} \mathbf{Pr}(\text{page fault} | X_t = \mathbf{s}) \cdot \mathbf{Pr}(X_t = \mathbf{s}).$$

In steady state, we can replace the above probabilities with their limiting values and obtain

$$F(H_{FIFO-LRU}(k)) = \sum_{\mathbf{s} \in \mathbf{Q}} (1 - \sum_{i=1}^m \beta_{j_i}) \cdot \pi_{\mathbf{s}} = \sum_{\mathbf{s} \in \mathbf{Q}} \left[G^{-1}(k) D_1^2(\mathbf{s}) \frac{\prod_{i=1}^m \beta_{j_i}}{\prod_{i=1}^{m-k+1} D_i(\mathbf{s})} \right] \quad (3)$$

as desired. □

Corollary 1: The steady-state fault rate for the pure LRU policy is given by

$$F(LRU) = \sum_{\mathbf{s} \in \mathbf{Q}} \left[D_1^2(\mathbf{s}) \frac{\prod_{i=1}^m \beta_{j_i}}{\prod_{i=1}^m D_i(\mathbf{s})} \right].$$

Proof: For $k=1$, the $H_{FIFO-LRU}$ policy replaces the same page that the pure LRU policy replaces. Thus, $F(H_{FIFO-LRU}(1))=F(LRU)$. Substituting $k=1$ into equation (3), we immediately have the desired result. □

Corollary 2: The steady-state fault rate for the pure FIFO policy is given by

$$F(FIFO) = \sum_{\mathbf{s} \in \mathbf{Q}} [G^{-1}(m) D_1(\mathbf{s}) \prod_{i=1}^m \beta_{j_i}].$$

Proof: For $k=m$ (i.e., $\mathbf{B}=\emptyset$), the $H_{FIFO-LRU}$ policy degenerates into the pure FIFO policy. Therefore, $F(H_{FIFO-LRU}(m))=F(FIFO)$. The result follows trivially by setting $k=m$ in equation (3). □

3.1.2. The RAND-LRU Hybrid Policy

In this section, we consider the simple variant of the $H_{FIFO-LRU}$ policy where the page to be moved from $\mathbf{T}$ to $\mathbf{B}$ at the time of a replacement is selected at random, uniformly over the pages that currently constitute $\mathbf{T}$. More precisely, the algorithm is identical to the $H_{FIFO-LRU}$ policy except that the top-to-bottom replacement is performed according to the RAND policy.

We proceed with the analysis after giving the problem a formulation identical to that in the previous section. For this policy, the one-step transition probabilities are

$$p(\mathbf{s}, \mathbf{s}') = \begin{cases} \sum_{i=1}^k \beta_{j_i}, & \text{if } \mathbf{s}' = \mathbf{s} \\ \beta_j/k, & \text{if } \mathbf{s}' = [j \ j_1 j_2 \dots j_{h-1} j_{h+1} \dots j_k j_h j_{k+1} \dots j_{m-1}] \\ & \text{where } j \notin \mathbf{s} \text{ and } 1 \leq h \leq k \\ \beta_l/k, & \text{if } \mathbf{s}' = [j_l j_1 j_2 \dots j_{h-1} j_{h+1} \dots j_k j_h j_{k+1} \dots j_{l-1} j_{l+1} \dots j_m] \\ & \text{where } 1 \leq h \leq k \text{ and } k < l \leq m \\ 0, & \text{otherwise.} \end{cases}$$

Note that all pages in $\mathbf{T}$ are eligible for replacement with the same probability $1/k$. It can easily be demonstrated that the resulting Markov chain is ergodic and therefore π exists and is a solution to equation (1) with the above one-step transition probabilities. For state $\mathbf{s} = [j_1, j_2, j_3, \dots, j_m]$, equation (1) can be written as:

$$\pi_{\mathbf{s}} = \pi_{\mathbf{s}} \sum_{i=1}^k \beta_{j_i} + \frac{\beta_{j_1}}{k} \left[\sum_{j \notin \mathbf{s}} \pi_{\mathbf{u}_j} + \sum_{h=1}^k \sum_{k < l \leq m} \pi_{\mathbf{v}_l} \right] \quad (4)$$

where $\mathbf{u}_j = [j_2, j_3, \dots, j_h, j_{k+1}, j_{h+1}, \dots, j_k, j_{k+2}, \dots, j_m, j]$ and $\mathbf{v}_l = [j_2, j_3, \dots, j_h, j_{k+1}, j_{h+1}, \dots, j_k, j_{k+2}, \dots, j_l, j_1, j_{l+1}, \dots, j_m]$ have again been constructed to result in state $\mathbf{s}$ upon the referencing of page j_1 .

Lemma 2: For the $H_{\text{RAND-LRU}}$ policy with parameter k , the equilibrium probability of state $\mathbf{s} = [j_1, j_2, j_3, \dots, j_m]$ is given by the same expression as for the $H_{\text{FIFO-LRU}}$ policy with the same parameter.

Proof: As in the proof of Lemma 1, we substitute the proposed solution for $\pi_{\mathbf{s}}$ into equation (4) and reduce it to an identity. □

Theorem 2: The steady-state fault rate for the $H_{\text{RAND-LRU}}$ policy with parameter k is equal to that of the $H_{\text{FIFO-LRU}}$ policy, i.e.,

$$F(H_{\text{RAND-LRU}}(k)) = F(H_{\text{FIFO-LRU}}(k))$$

Proof: The proof trivially follows from Lemma 2 by conditioning on the states of the Markov chain.

□

Corollary 3: For programs whose behavior is perfectly represented by the IRM, the FIFO and RAND policies result in identical steady-state fault rates. That is,

$$F(FIFO) = F(RAND) \quad \text{under the IRM.}$$

Proof. The proof trivially follows from Theorem 2 upon observing that, for $k=m$, the $H_{RAND-LRU}$ policy degenerates into the pure RAND policy and the $H_{FIFO-LRU}$ policy degenerates into the pure FIFO policy.

□

Note that the above result has been obtained through a different method by Gelenbe [16].

3.2. Variable Partition Hybrids

In this section we consider hybrid policies that use WS management for $\mathbf{B}$, thus resulting in variable size memory partitions. Note, however, that the partition size can never become less than k pages (the size of $\mathbf{T}$), where k is a static parameter of the policy. Recall that the pure WS algorithm with parameter τ retains a page in memory only if it has been referenced at least once during the previous τ time units [10]. In our case, since we have no information about references to a page during the period of its membership in $\mathbf{T}$, we must somehow estimate the last time the page was referenced when it leaves $\mathbf{T}$.

3.2.1. The FIFO-WS Hybrid Policy

If r_t is the page referenced at instant t , we define the following terms:

- (i) An *event* is said to occur at time t if $r_t \notin \mathbf{T}$.
- (ii) A *page fault* is said to occur at time t if $r_t \notin \{\text{memory}\}$.

Obviously, all page faults are also events (since $\{\text{memory}\} = \mathbf{T} \cup \mathbf{B}$). Furthermore, with each page $i \in \mathbf{B}$, we associate a time t_i that is an estimate of the time of last reference to that page.

We can now describe the operation of the $H_{FIFO-WS}$ policy with parameters k and τ .

W1: (Null case) If $r_t \in \mathbf{T}$, no action is taken (as before, no action *can* be taken).

W2: (Event or fault) If $r_t \notin \mathbf{T}$, then $\mathbf{T} \leftarrow \mathbf{T} + r_t - i$, where page i is the FIFO page in $\mathbf{T}$. As before, “+” and “-” denote set operations. We provide our estimate of the time of last reference to page i as the current time: $t_i \leftarrow t$. We update $\mathbf{B}$ by removing all pages with last reference time estimates earlier than $t - \tau$; $\mathbf{B} \leftarrow \mathbf{B} - \mathbf{J}$, where $\mathbf{J} = \{j \in \mathbf{B}; t_j \leq t - \tau\}$.

Note that, in the above, we know page i could not have been referenced at time t since at that instant an event occurred and page i was in $\mathbf{T}$. Alternative estimates for the time of last reference to a page during its membership in $\mathbf{T}$ are the time of entry into $\mathbf{T}$ and the arithmetic mean of the times of entry and exit from $\mathbf{T}$. In practice, however, these variations have a negligible effect on performance [3].

For the variable partition hybrid policies, we are interested in obtaining expressions not only for the steady-state fault rate, $F(A)$, but also for the mean memory occupancy, $M(A)$, which is the expected number of pages that are in memory in steady state.

We proceed with the analysis after the following definitions and preliminaries. Let I_i be a random variable such that

$$I_i = \begin{cases} 1, & \text{if page } i \text{ is in memory} \\ 0, & \text{otherwise.} \end{cases}$$

Let the random variable ψ_i denote the number of events that have occurred since page i last left $\mathbf{T}$. As before, r_t represents the page name corresponding to the current reference. Note that, if $\mathbf{T}$ is examined in isolation, references that cause events will always correspond to page faults in a memory of k page frames that is managed by the pure FIFO policy. Consider a Markov chain formulation of the FIFO policy completely analogous to our formulation in section 3.1.1 [18, 7]. Let $\mathbf{Q}$ be the state space consisting of the states $\mathbf{s} = [j_1, j_2, j_3, \dots, j_k]$, representing all combinations of n items taken k at a time (as it turns out, for the FIFO and RAND analysis, knowledge of the ordering of pages within a combination is not necessary). For a given page i , partition the state

space into two disjoint sets $\mathbf{R}_i$ and $\mathbf{P}_i$ such that $\mathbf{R}_i$ contains all the states that include page i and $\mathbf{P}_i$ all other states. In other words, let $\mathbf{R}_i \subset \mathbf{Q}$ be such that $\mathbf{s} \in \mathbf{R}_i$ iff $i \in \mathbf{s}$ where $\mathbf{R}_i \cup \mathbf{P}_i = \mathbf{Q}$ and $\mathbf{R}_i \cap \mathbf{P}_i = \emptyset$. Let π_i denote the steady-state probability that page i is an element of $\mathbf{T}$. Then, from our pure FIFO analysis, the steady-state occupancy probabilities for the states of a memory of k page frames are known. The desired expression is obtained simply by summing the limiting state occupancy probabilities over all those states that contain page i . That is,

$$\pi_i = \Pr(i \in \mathbf{T}) = \sum_{\mathbf{s} \in \mathbf{R}_i} \pi_{\mathbf{s}} = G^{-1}(k) \sum_{\mathbf{s} \in \mathbf{R}_i} \prod_{j \in \mathbf{s}} \beta_j.$$

Let p_i denote the steady-state probability of an event given that page i is not in $\mathbf{T}$. That is,

$$p_i = \lim_{t \rightarrow \infty} \Pr(r_t \notin \mathbf{T} \mid i \notin \mathbf{T}).$$

Again from our Markov chain analysis of the pure FIFO policy, we obtain

$$p_i = \lim_{t \rightarrow \infty} \frac{\Pr(r_t \notin \mathbf{T}, i \notin \mathbf{T})}{\Pr(i \notin \mathbf{T})} = \frac{\sum_{\mathbf{s} \in \mathbf{P}_i} [(1 - \sum_{j \in \mathbf{s}} \beta_j) \prod_{j \in \mathbf{s}} \beta_j]}{G(k) (1 - \pi_i)}.$$

Let Γ_i be the steady-state probability that page i is referenced given that an event has occurred and page i is not in $\mathbf{T}$. Using basic probability concepts and earlier definitions, Γ_i can be expressed as

$$\begin{aligned} \Gamma_i &= \lim_{t \rightarrow \infty} \Pr(r_t = i \mid r_t \notin \mathbf{T}, i \notin \mathbf{T}) \\ &= \frac{(1 - \pi_i) \beta_i}{\sum_{\mathbf{s} \in \mathbf{P}_i} D_1(\mathbf{s}) \pi_{\mathbf{s}}}. \end{aligned}$$

Given the above background, we can prove the following:

Lemma 3: The steady-state probability that page i is in a memory that is managed by the $H_{FIFO-WS}$ policy with parameters k and τ is given by

$$\Pr(I_i = 1) = \pi_i + (1 - \pi_i) \cdot \sum_{j=1}^{\tau} \left[\Gamma_i (1 - \Gamma_i)^{j-1} \cdot \sum_{l=j}^{\tau} \binom{l-1}{l-j} p_i^j (1 - p_i)^{l-j} \right] \quad (5)$$

where

$$\pi_i = \Pr(i \in \mathbf{T}) = \sum_{\mathbf{s} \in \mathbf{R}_i} \pi_{\mathbf{s}},$$

$$\Gamma_i = \lim_{l \rightarrow \infty} \Pr(r_l = i \mid r_l \notin \mathbf{T}, i \notin \mathbf{T}),$$

and

$$p_i = \lim_{l \rightarrow \infty} \Pr(r_l \notin \mathbf{T} \mid i \notin \mathbf{T}).$$

The proof is presented in the Appendix.

Theorem 3: The steady-state fault rate for the $H_{FIFO-WS}$ policy with parameters k and τ is given by

$$F(H_{FIFO-WS}) = 1 - \sum_{i=1}^n \beta_i \left[\pi_i + (1 - \pi_i) \cdot \sum_{j=1}^{\tau} \left[\Gamma_i (1 - \Gamma_i)^{j-1} \sum_{l=j}^{\tau} \binom{l-1}{l-j} p_i^j (1 - p_i)^{l-j} \right] \right],$$

where π_i , Γ_i , and p_i are as defined in Lemma 3.

Proof: In steady state, the probability that page i causes a fault is simply the steady-state probability that that page is not in memory and is referenced. Thus, conditioning on the page,

$$F(H_{FIFO-WS}) = \sum_{i=1}^n (1 - \Pr(I_i = 1)) \cdot \beta_i.$$

The result then follows trivially upon the substitution of equation (5) into the above expression. □

Theorem 4: The mean memory occupancy due to the $H_{FIFO-WS}$ policy with parameters k and τ is given by

$$M(H_{FIFO-WS}) = \sum_{i=1}^n \left[\pi_i + (1 - \pi_i) \cdot \sum_{j=1}^{\tau} \left[\Gamma_i (1 - \Gamma_i)^{j-1} \sum_{l=j}^{\tau} \binom{l-1}{l-j} p_i^j (1 - p_i)^{l-j} \right] \right]$$

where π_i , Γ_i , and p_i are as defined in Lemma 3.

Proof: Conditioning on the page,

$$M(H_{FIFO-WS}) = \sum_{i=1}^n \Pr(I_i=1).$$

The result follows immediately from the substitution of equation (5) into the above expression. □

3.2.2. The RAND-WS Hybrid Policy

Here we consider a variation of the $H_{FIFO-WS}$ policy that operates exactly as described in the previous section except that, at the times of events, the page to leave the top is selected at random uniformly over the pages currently in $\mathbf{T}$.

Theorem 5: The steady-state fault rates and the mean memory occupancies for the $H_{RAND-WS}$ policy are identical to those of the $H_{FIFO-WS}$ policy with the same parameters.

Proof. By Lemma 2 we know that the equilibrium probabilities for the Markov chain states are the same for both the RAND and the FIFO policies. Also, by Corollary 3, we know that the two policies result in the same steady-state fault rate. Based on these observations, the proof follows trivially. □

4. Numerical Results

In this and subsequent sections, we restrict our attention to the fixed partition hybrid policies. Parallel results for the variable partition hybrid policies can be found in [3]. Given the closed-form expression for $F(H_{FIFO-LRU}(k))$ in equation (3) (recall that $F(H_{RAND-LRU}(k))$ has the same expression), we are interested in its functional dependency on the policy parameter k for various values of m and of the IRM parameters. However, due to the complexity of the expression, its form between the two end points $k=1$ and $k=m$ (which correspond to pure LRU and pure FIFO or pure RAND, respectively) is difficult to study analytically. We know that, for the IRM, the fault rate due to pure LRU is always less than or equal to the FIFO/RAND fault rate

Figure 1. $H_{FIFO-LRU}$ and $H_{RAND-LRU}$ fault rates for two sample programs as a function of the policy parameter k .

for all m and for all model parameters [18].

To study the behavior of the $F(H_{FIFO-LRU}(k))$ function between the two end points we computed numerical values of the expression for various instances of the program model. To minimize the number of parameters involved, the two instances of the IRM we consider were generated through the equations $\beta_i = ci$ and $\beta_i = c^i$, which are called the *arithmetic* and the *geometric* models, respectively. In both cases, the constant c is chosen such that $\sum_{i=1}^n \beta_i = 1$. In Figure 1, $F(H_{FIFO-LRU}(k))$ is plotted as a function of k for the fixed values $n=8$ and $m=7$. Note the strong convexity of the two curves; particularly the one corresponding to the geometric IRM, where fault rates that are practically identical to the pure LRU fault rate are achieved even for top sizes of 5 pages (or equivalently, bottom sizes of 2 pages). Note that, although not indicated in Figure 1, for the case where $\beta_i = 1/n$ (uniform distribution of the IRM parameters), *all* demand algorithms result in the same fault rate of $(n-m)/n$. For this case, $F(H_{FIFO-LRU}(k))$ obviously does not vary with k . Outside of this degenerate case, the strict convexity of $F(H_{FIFO-LRU}(k))$ as a function k for all m and IRM parameter values remains a conjecture.

The above results have to be interpreted with caution for two reasons:

- (i) They are based on a model of program behavior that is known to lack many of the properties of real programs,
- (ii) Due to the factorial growth of the complexity of the expressions involved, numerical results can only be presented for unrealistically small values of the program size, n , and of the memory size, m .

To eliminate these concerns, trace-driven simulations of the hybrid policies were carried out using three program samples. The reader should consult [3] for further details about the programs and the simulation procedure. The resulting values of the fault rate as a function of the policy parameter are displayed in Figure 2. These curves indicate that the conclusions based on analytic results are also valid for the three programs available to us. All three programs display a strongly convex curve with a well-defined knee at approximately $k=40$ pages. Note that this represents a

(a)

(b)

(c)

Figure 2. $H_{FIFO-LRU}$ and $H_{RAND-LRU}$ fault rates as a function of the policy parameter k for (a) WATFIV, (b) APL and (c) FFT.

bottom size of 10 pages or, equivalently, 20 percent of the total memory size. We also note that the $H_{FIFO-LRU}$ and $H_{RAND-LRU}$ policies result in different fault rates for these programs. This confirms the fact that these programs do not satisfy the IRM assumptions. No general conclusion about the relative performances of the $H_{FIFO-LRU}$ and $H_{RAND-LRU}$ policies, however, can be derived from these results, as each is uniformly superior for one of the programs while both perform about the same for the remaining program.

5. Cost Considerations

Up to this point, the only performance measure we have been concerned with has been the steady-state fault rate as a function of the mean memory occupancy. The operation of these hybrid policies requires setting the policy parameter k to some value. If the cost of referencing a page that is in $\mathbf{B}$ were negligible, then setting $k=1$ would almost always produce optimum performance with respect to the page fault rate (there are few points in the graphs of Figure 2 where the fault rate actually drops as k is increased beyond 1). However, since a finite cost is incurred each time a page in $\mathbf{B}$ is referenced, the selection of the policy parameter should be guided by the desire to keep the fault rate close to the value given by pure LRU (i.e., corresponding to $k=1$) while minimizing the size of the set $\mathbf{B}$ (i.e., minimizing the rate of reclaims). Intuitively, the policy should be operated with the parameter set to a value close to the knee that occurs in all three fault rate graphs. Formally, we define a performance measure, $C(\cdot)$, that is the weighted sum of the fault rate and the reclaim rate for a given value of the policy parameter. For a page replacement algorithm A and a ratio of page fault service time to page reclaim service time given by α , let

$$C(A, m, k, \alpha) = f(A, m-k) + \alpha \cdot F(A, m, k),$$

where m is the memory size, k is the policy parameter (or the size of $\mathbf{T}$) and $f(\cdot)$ is the reclaim rate. Note that, while appropriate as a *responsiveness* measure from the point of view of an individual program, the cost function $C(\cdot)$ is inappropriate for multiprogramming system *throughput* considerations since the major part of the delay due to a page fault results from the paging I/O

operation, which can be overlapped with other CPU activity. On the other hand, a reclaim operation is performed by the CPU and cannot be overlapped.

Due to the complexity of the expressions for the steady-state fault and reclaim rates, an analytic minimization of the cost function defined with respect to k cannot be carried out. As before, we resort to trace-driven simulation to study the behavior of this cost function as the policy parameter changes. Figure 3 displays this cost function for the three programs under the $H_{FIFO-LRU}$ policy. For the sake of brevity we restrict the presentation of the results to the case $\alpha=100$ under the $H_{FIFO-LRU}$ policy. Our conclusions, however, are applicable to the $H_{RAND-LRU}$ policy as well as to the cases when $\alpha=10$ and $\alpha=1000$ under both policies. The particular choice of the value of α for the results reported here is further supported by measurements from an actual system [4]. As can be seen in Figure 3, the optimum setting for the policy parameter k varies significantly depending on the program as well as on the total amount of memory allocated to it. In [24], Turner and Levy report on the behavior of this cost function for different values of α . The observed insensitivity of the optimum value of the policy parameter to variations of α is unfortunately not very valuable since, for a given implementation, α tends to be rather static anyway. In a real system, the characteristics of the different programs and the amount of memory allocated to each program are far more dynamic. Our results demonstrate that it is precisely these variations that require the parameter to be adjusted for the hybrid policies to perform optimally. Conditions which make m a variable rather than a fixed system parameter are discussed in the next section.

6. Extensions

In a multiprogramming environment, the fixed partition hybrids can be extended in a natural way to operate as FIFO-Global LRU and RAND-Global LRU hybrids, thus resulting in variable partitions for the individual programs. This can be accomplished simply by maintaining a single fixed size bottom containing the reclaimable pages of all the programs that are currently being multiprogrammed (the tops for each of the programs, however, are still maintained

(a)

(b)

(c)

Figure 3. The weighted sum cost observed under the $H_{FIFO-LRU}$ policy for varying amounts of memory as a function of the policy parameter and programs (a) WATFIV, (b) APL and (c) FFT.

separately). In such an environment, the total number of page frames allocated to a program at any given time will be the size of its top plus a random variable that represents the number of pages in $\mathbf{B}$ belonging to the given program. Under this extension, the study of the performance of individual programs is severely complicated due to their interactions with the other concurrently running programs.

Although we have not studied it analytically, we note that this extension of the $H_{FIFO-LRU}$ policy adequately models the memory management policy employed in the VMS operating system for the VAX-11/780 computer system [9, 24, 19]. Unfortunately, for optimality to be preserved under this extension, not only must the hybrid policy parameter be set on a per-program basis, but it should also vary dynamically as the amount of memory allocated to the program changes. We note that suitable heuristics that try to keep the top size to the bottom size ratio constant or that try to maintain a constant reclaim rate can be employed to vary the policy parameter dynamically. The current version of VMS, in fact, implements a scheme similar to the second option to adjust the size of $\mathbf{T}$ as each program executes [24].

7. Conclusions

We have introduced a class of hybrid algorithms that are suitable for page replacement decisions in hierarchical stores that lack hardware support other than the address translation mechanism. Expressions for the steady-state fault rates generated by these policies have been derived based on the Independent Reference Model of program behavior. With the appropriate setting of the policy parameter, numerical and empirical results suggest that these algorithms are capable of achieving fault rates close to those of the pure LRU and WS policies while incurring costs comparable to those of the FIFO and RAND policies.

We have defined a cost function that includes explicitly the cost of reclaim operations; this function is more suitable than the simple fault rate for the evaluation of hybrid algorithms. Based on this cost function, the parameter of the hybrid policies we analyzed was seen to be sensitive to variations in the program characteristics as well as to the amount of memory allocated to

it. We have suggested reasonable heuristics that can be incorporated into these policies to allow the dynamic adjustment of their parameters.

APPENDIX

Proof of Lemma 1: It suffices to show that the proposed solution satisfies equation (2). First, note that, from its definition, the $D_i()$ function for the states $\mathbf{u}_j$ and $\mathbf{v}_l$ can be written as

$$D_i(\mathbf{u}_j) = 1 - \sum_{l=2}^{m-i+2} \beta_{j_l} = D_{i-1}(\mathbf{s}) + \beta_{j_1}, \quad i > 1 \quad (\text{A1})$$

and

$$D_i(\mathbf{v}_l) = \begin{cases} D_i(\mathbf{s}) & , \quad i \leq m-l+1 \\ D_{i-1}(\mathbf{s}) + \beta_{j_1} & , \quad i > m-l+1 \end{cases} \quad (\text{A2})$$

In terms of $D_i(\mathbf{s})$, equation (2) becomes

$$\pi_{\mathbf{s}} = \pi_{\mathbf{s}} (1 - D_{m-k+1}(\mathbf{s})) + \beta_{j_1} \left[\sum_{j \notin \mathbf{s}} \pi_{\mathbf{u}_j} + \sum_{k < l \leq m} \pi_{\mathbf{v}_l} \right].$$

Substituting the proposed solution into the above equation, we have

$$\pi_{\mathbf{s}} D_{m-k+1}(\mathbf{s}) = \beta_{j_1} \left[\sum_{j \notin \mathbf{s}} \frac{\beta_j \prod_{i=2}^m \beta_{j_i}}{G(k) \prod_{i=2}^{m-k+1} D_i(\mathbf{u}_j)} + \sum_{k < l \leq m} \frac{\prod_{i=1}^m \beta_{j_i}}{G(k) \prod_{i=2}^{m-k+1} D_i(\mathbf{v}_l)} \right].$$

Using equations (A1) and (A2) for $D_i(\mathbf{u}_j)$ and $D_i(\mathbf{v}_l)$, respectively, in the above equation and simplifying, we obtain

$$\pi_{\mathbf{s}} \frac{G(k) D_{m-k+1}(\mathbf{s})}{\prod_{i=1}^m \beta_{j_i}} = \frac{D_i(\mathbf{s})}{\prod_{i=1}^{m-k} (D_i(\mathbf{s}) + \beta_{j_1})} + \sum_{k < l \leq m} \frac{\beta_{j_1}}{\prod_{i=2}^{m-l+1} D_i(\mathbf{s}) \prod_{i=m-l+1}^{m-k} (D_i(\mathbf{s}) + \beta_{j_1})}, \quad (\text{A3})$$

where we have used the fact that $\sum_{j \notin \mathbf{s}} \beta_j = D_1(\mathbf{s})$.

Putting the right hand side terms over a common denominator and applying the transformations $y=m-k$ and $z=l-k$ to the indices results in

$$\pi_s \frac{G(k) \prod_{i=2}^{m-k+1} D_i(\mathbf{s})}{\prod_{i=1}^m \beta_{j_i}} = \frac{\prod_{i=1}^y D_i(\mathbf{s}) + \beta_{j_1} \cdot \sum_{1 \leq z \leq y} \left[\prod_{i=y-z+2}^y D_i(\mathbf{s}) \cdot \prod_{i=1}^{y-z} (D_i(\mathbf{s}) + \beta_{j_1}) \right]}{\prod_{i=1}^y (D_i(\mathbf{s}) + \beta_{j_1})} . \quad (\text{A4})$$

Through the identity

$$\prod_{i=1}^a (Z_i + \lambda) = \prod_{i=1}^a Z_i + \lambda \cdot \sum_{1 \leq b \leq a} \left[\prod_{i=1}^{a-b} (Z_i + \lambda) \cdot \prod_{i=a-b+2}^a Z_i \right],$$

the right hand side of equation (A4) reduces to unity, and finally we find that

$$\pi_s = \frac{\prod_{i=1}^m \beta_{j_i}}{G(k) \prod_{i=2}^{m-k+1} D_i(\mathbf{s})} = \pi_s$$

as claimed. □

Proof of Lemma 3: Conditioning on the position of page i in memory, we obtain

$$\begin{aligned} \Pr(I_i=1) &= \Pr(I_i=1 \mid i \in \mathbf{T}) \cdot \Pr(i \in \mathbf{T}) + \Pr(I_i=1 \mid i \notin \mathbf{T}) \cdot \Pr(i \notin \mathbf{T}) \\ &= 1 \cdot \pi_i + (1 - \pi_i) \cdot \Pr(I_i=1 \mid i \notin \mathbf{T}), \end{aligned} \quad (\text{A5})$$

since pages in $\mathbf{T}$ are always in memory. We now derive an expression for the probability that a page is still in memory given that it is not in $\mathbf{T}$. We proceed by conditioning on the number of events that have occurred since page i last left $\mathbf{T}$:

$$\Pr(I_i=1 \mid i \notin \mathbf{T}) = \sum_{j=1}^{\infty} \Pr(I_i=1 \mid i \notin \mathbf{T}, \psi_i=j) \cdot \Pr(\psi_i=j \mid i \notin \mathbf{T}) . \quad (\text{A6})$$

Now, based on the IRM assumption, given that page i is not in $\mathbf{T}$, referencing it at the instant of events constitutes independent Bernoulli trials [13] with probability of success given by Γ_i . Therefore, the number of events that have occurred since page i last left $\mathbf{T}$ has a geometric distribution with parameter Γ_i . That is,

$$\Pr(\psi_i=j \mid i \notin \mathbf{T}) = \Gamma_i(1-\Gamma_i)^{j-1}, \quad j \geq 1. \quad (\text{A7})$$

Returning to equation (A6), the probability that page i is in memory given that it is not in $\mathbf{T}$ and j events have occurred since it left $\mathbf{T}$ can be expressed as

$$\Pr(I_i=1 \mid i \notin \mathbf{T}, \psi_i=j) = \Pr(T_1 + T_2 + \dots + T_j \leq \tau),$$

where the T_i 's are the interevent times given that page i is not in $\mathbf{T}$ and τ is the policy parameter. Recall that a reference to page i causes an event with probability p_i given that $i \notin \mathbf{T}$. Thus, with respect to events, memory references constitute independent Bernoulli trials with probability of success p_i under the given condition. Hence, the number of references between events has a geometric distribution with parameter p_i . That is,

$$\Pr(T_i=x) = p_i(1-p_i)^{x-1}, \quad x \geq 1.$$

Finally, the random variable $T_1 + T_2 + \dots + T_j$, which is the sum of j independent geometric random variables, has a negative binomial distribution with parameters p_i and j [13]. In other words,

$$\Pr(T_1 + T_2 + \dots + T_j \leq x) = \sum_{l=j}^x \binom{l-1}{j-1} p_i^j (1-p_i)^{l-j}, \quad x \geq j. \quad (\text{A8})$$

Substituting equations (A7) and (A8) into equation (A6), we obtain

$$\Pr(I_i=1 \mid i \notin \mathbf{T}) = \sum_{i=1}^{\tau} \left[\Gamma_i(1-\Gamma_i)^{j-1} \sum_{l=j}^{\tau} \binom{l-1}{j-1} p_i^j (1-p_i)^{l-j} \right], \quad (\text{A9})$$

where the first summation terminates at τ since $\Pr(T_1 + T_2 + \dots + T_j \leq \tau) = 0$ for all $j > \tau$ due to the fact that $T_i \geq 1$. Combining equations (A9) and (A5), the desired result is obtained. □

8. References

- [1] A. V. Aho, P. J. Denning, and J. D. Ullman, "Principles of Optimal Page Replacement," *J. ACM* **18** pp. 80-93 (January 1971).

- [2] O. I. Aven, L. B. Boguslavsky, and Y. A. Kogan, "Some Results on Distribution-Free Analysis of Paging Algorithms," *IEEE Trans. Comptrs.* **C-25(7)** pp. 737-745 (July 1976).
- [3] Ö. Babaoğlu, "Virtual Storage Management in the Absence of Reference Bits," Ph.D. Th., Univ. of California, Berkeley, Memorandum No. ERL M81/92 (November 1981).
- [4] Ö. Babaoğlu and W. N. Joy, "Converting a Swap-Based System to do Paging in an Architecture Lacking Page-Referenced Bits," *Proc. Eighth Symp. Operating Systems Principles, SIGOPS*, pp. 78-86 (December 1981).
- [5] F. Baskett and A. Rafii, "The A_0 Inversion Model of Program Paging Behavior," Stanford U. Comp. Sci. Dept., Report STAN-CS-76-579 (October 1976).
- [6] L. A. Belady, "A Study of Replacement Algorithms for a Virtual Storage Computer," *IBM Syst. J.* **5** pp. 78-101 (1966).
- [7] E. G. Coffman and P.J. Denning, *Operating Systems Theory*, Prentice-Hall, Englewood Cliff, New Jersey (1973).
- [8] F. J. Corbato, "A Paging Experiment with the Multics System," Project MAC Memo MAC-M-384 Mass. Inst. of Tech. (July 1968). Published in In Honor of P. M. Morse ed. Ingard MIT Press 1969, pp. 217-228.
- [9] DEC, *VAX-11/780 Software Handbook*, Digital Press (1978).
- [10] P. J. Denning, "The Working Set Model of Program Behavior," *Comm. ACM* **11(5)** pp. 323-333 (May 1968).
- [11] Peter J. Denning, "Virtual Memory," *Computng. Surveys* **2(3)** pp. 153-189 (September 1970).
- [12] M. Easton and P. A. Franaszek, "Use Bit Scanning in Replacement Decisions," *IEEE Trans. Comptrs.* **C-28** pp. 133-141 (February 1979).
- [13] W. Feller, *Introduction to Probability and its Applications*, Wiley (1968). 3rd. ed., vol. 1.

- [14] M. H. Fogel, "The VMOS Paging Algorithm - A Practical Implementation of a Working Set Model," *Operating Sysys. Rev.* **8** pp. 8-17 (January 1974).
- [15] John Fotheringham, "Dynamic Storage Allocation in the Atlas Computer Including an Automatic Use of Backing Store," *Comm. ACM* **4**(10) pp. 435-436 (October 1961).
- [16] E. Gelenbe, "A Unified Approach to the Evaluation of a Class of Replacement Algorithms," *IEEE Trans. Comptrs.* **C-22** pp. 611-618 (June 1973).
- [17] T. Kilburn, D. B. G. Edwards, M. J. Lanigan, and F. H. Sumner, "One Level Storage Systems," *IRE Trans. Electronic Computers* **EC-11** pp. 223-235 (April 1962).
- [18] W. F. King, "Analysis of Demand Paging Algorithms," *Proc. IFIPS Congress*, pp. TA-3-155 - TA-3-159 , Ljubljana, Yugoslavia(1971).
- [19] Edward D. Lazowska, "The Benchmarking, Tuning and Analytical Modeling of VAX/VMS," *Proceedings of the Conference on Simulation, Measurement and Modeling of Computer Systems, Boulder, Colorado*, pp. 57-64 (August 1979).
- [20] R. L. Mattson, J. Gecsei, D. R. Slutz, and I. L. Traiger, "Evaluation Techniques for Storage Hierarchies," *IBM Syst. J.* **9** pp. 78-117 (1970).
- [21] B. G. Prieve, "A Page Partition Replacement Algorithm," Ph.D. Th., Univ. of California Berkeley, California (1974).
- [22] S. M. Ross, *Applied Probability Models with Optimization Applications*, Holden-Day, San Francisco (1970).
- [23] A. J. Smith, "Bibliography on Paging and Related Topics," *Operating Sysys. Rev.* **12**(4) pp. 39-56 (October 1978).
- [24] R. Turner and H. Levy, "Segmented FIFO Page Replacement," *Performance Eval. Rev.* **10**(3) pp. 48-51 (Fall 1981). Proceeding of ACM/SIGMETRICS Coference on Measurement and Modeling of Computer Systems, Las Vegas, Nevada.