

---

# Generalized Working Sets for Segment Reference Strings

Peter J. Denning  
Purdue University

Donald R. Slutz  
IBM, San Jose

---

The working-set concept is extended for programs that reference segments of different sizes. The generalized working-set policy (GWS) keeps as its resident set those segments whose retention costs do not exceed their retrieval costs. The GWS is a model for the entire class of demand-fetching memory policies that satisfy a resident-set inclusion property. A generalized optimal policy (GOPT) is also defined; at its operating points it minimizes aggregated retention and swapping costs. Special cases of the cost structure allow GWS and GOPT to simulate any known stack algorithm, the working set, and VMIN. Efficient procedures for computing demand curves showing swapping load as a function of memory usage are developed for GWS and GOPT policies. Empirical data from an actual system are included.

**Key Words and Phrases:** database referencing, memory management, optimal memory policies, paging, program behavior, program measurement, segmentation, working sets

**CR Categories:** 4.30, 8.10

---

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

This work was supported at Purdue University in part by NSF grant GJ-41289.

Authors' addresses: Peter J. Denning, Computer Sciences Department, Purdue University, West Lafayette, IN 47907. Donald R. Slutz, IBM Research Laboratory, 5600 Cottle Road, San Jose, CA 95193.

© 1978 ACM 0001-0782/78/0900-0750 \$00.75

## Introduction

Segment referencing has become increasingly prevalent. It is used in many database systems; it is supported in hardware by several manufacturers; it underlies "program restructuring," which seeks a locality-preserving distribution of small logical program blocks among large physical storage pages. There is a clear need for memory demand measures under segment referencing.

The common procedures for measuring memory demand are suited only for paged memory systems. They measure a resident set's size by counting pages, and the swapping load by counting page faults. But if storage blocks are segments of different sizes, these measures do not accurately portray a program's memory demands. The number of resident segments may bear little relation to the memory required to hold them; the number of missing-segment faults may not measure the load actually placed on the swapping system.

This paper presents the *generalized working set* (GWS) approach to measuring memory demand under segment referencing. It is based on defining the cost of retaining a segment in memory without being referenced, and the cost of retrieving (swapping in) a missing segment. The GWS memory policy assigns each active program a resident set containing each segment whose retention cost does not exceed its retrieval cost. A parameter  $\Theta$ , the *threshold*, represents the cost of retrieval relative to retention. The GWS models all one-parameter memory policies whose resident sets satisfy an inclusion property under increasing values of the control parameter ( $\Theta$ ). The well-known "stack algorithms" [20, 6] and "time window working set" [9, 14, 6] are special cases of the GWS. This model extends the measurement technique to segment referencing, and it unifies previous models as well.

In a single pass over a given segment-reference string, GWS procedures measure a *memory demand curve*  $y = f(x)$ . Each possible threshold value ( $\Theta$ ) generates a particular demand point  $(x, y)$  on this curve. This curve shows the tradeoff between a "memory-space investment,"  $x$ , and a "swapping load"  $y$ . The memory-space investment  $x$  is either  $\sigma$ , the mean size of the resident set, or  $s$ , the mean of the retention costs of segments kept in the GWS. The swapping load represents the delay from moving segments into main memory; it is represented as  $y = m + A\mu$ , where  $m$  is the missing-segment fault rate,  $\mu$  is the rate of information flow resulting from segment faults, and  $A$  is a parameter selected by the analyst. The familiar page-fault curve is the demand curve  $m = f(\sigma)$ . These definitions give the analyst considerable flexibility in choosing a memory-demand measure.

A special case of GWS policy is a *generalized optimal policy* (GOPT). No memory policy can generate a demand point below the demand curve of GOPT. Although GOPT's lookahead prevents it from operating in real time, its demand curves are easily obtained—in fact, the GOPT and GWS demand curves can be computed from each

other. It is thus cheap to learn how far from optimal a given GWS policy is. The GOPT policy reduces to VMIN for paging [21].

The original procedures for measuring page fault-rate curves under the time-window working-set policy required storage of order  $O(M + p)$ , where  $M$  is the maximum time window of interest and  $p$  is the number of pages. (See [14, 25].) In practice,  $M$  must be very large— $10^4$  or  $10^5$  references—to obtain demand points over the entire range of interest. The GWS and GOPT measurement procedures calculate demand points for  $N$  selected threshold values, with storage of order  $O(N + p)$ , where  $N$  can be as small as  $\log_2 M$ . For practical programs, this represents a storage reduction of two or more orders of magnitude and corresponding speedup in computing the demand curves. These procedures are generalizations of those noted by Easton and Bennett for the time-window working set [16], but they were developed independently by the authors of [11] and [24].

Because its cost functions measure each program singly, GWS analysis does not calculate the actual cost of running a program. It does not account for the overheads of placement or replacement policies, or the effects of queueing. GWS analysis does measure the tradeoff between memory-space investment and swapping load intrinsic to each given program.

To estimate the actual cost of running a program, the memory-demand curves of programs in a workload must be used to drive a simulation or analytic model of the entire system. The system model accounts for overhead and queueing. This has been done successfully many times with paging systems. For example, paging curves have been used to estimate processor utilization, throughput, and mean response time of various levels of multiprogramming [1, 2, 8, 12]; to determine bounds on throughput [13]; to construct synthetic workloads [27]; and to measure program locality [15, 18]. With demand curves from GWS analyses, these same techniques can be applied to systems with segment addressing.

## Demand Curves

A program's address space consists of  $p$  segments, denoted by indices  $1, \dots, p$ . The size of segment  $i$  is  $z_i$  data units. A data unit is a fixed quantity such as a bit, byte, word, or page. The total of segment sizes is  $Z = z_1 + \dots + z_p$ . Under paging, all  $z_i = 1$ .

A segment reference string is a sequence  $\{r(t), t = 1, \dots, T\}$  in which  $r(t)$  is the index number of the segment referenced at virtual time  $t$ . The *total volume of referenced information*,  $V$ , is the sum, over all  $t$ , of the size of the segment referenced at time  $t$ . Since we assume that the whole of a segment must be loaded in main memory for referencing, the mean resident set size of any memory policy is at least  $V/T$ .

Our analysis supposes that every program starts execution with an empty resident set, and that missing

segments are placed in the resident set on demand. A *memory policy* (MP) determines which segments are removed from resident sets. The MP's of interest here decide whether or not to retain or remove a segment by comparing *memory usage costs* against *swapping loads*. A fixed *threshold*,  $\Theta$ , specifies the relative cost the MP assigns to retaining and swapping segments. For a given such MP, each resident set is determined completely by the reference string and the setting of the control parameter,  $\Theta$ .

Our MP's associate a reference cost and a retention cost with each segment at each time  $t$ . The *reference cost* accounts for the unavoidable cost of using memory while a segment is being referenced. We let  $q_i$  denote the reference cost for segment  $i$ . The *total reference cost*,  $Q$ , is the sum, over all  $t$ , of the reference cost of the segment referenced at time  $t$ . The cost  $Q$  is incurred by every MP in processing the given reference string. For paging, all  $q_i = 1$ .

The (accumulated) *retention cost* accounts for the memory used to maintain a segment in residence beyond its prior reference. For each segment at time  $t$ , this cost is represented by a function  $R(i, t) \geq 0$  satisfying three properties:

- (a) Prior to the first reference of segment  $i$ ,  $R(i, t)$  is infinite: No finite expenditure can cause a segment to be resident before its first reference under a demand MP. (However, nonempty initial resident sets can be represented with  $R(i, 0)$  being a suitable finite value.)
- (b) If  $r(t) \neq i$ ,  $R(i, t + 1) \geq R(i, t)$ : Retention cost accumulates with time since prior reference.
- (c) If  $r(t) = i$ ,  $R(i, t^+) = 0$ : Retention cost is reset just after a reference. (The cost of the reference itself is accounted for by  $q_i$ .) In general, retention cost depends on some total internal state of (a model of) the program. Thus  $R(i, t)$  is *not* independent of  $R(j, t)$ . To keep the notation simple, we have not shown such a state explicitly as a parameter.

It is convenient to extend these definitions to continuous time, in which segments are referenced at integer times. In this case, segment  $i$  is regarded as being resident during  $[t, t + 1)$  whenever  $r(t) = i$ , and the cost of this reference is represented by  $q_i$ . (Note that  $R(i, t)$  need not be continuous.)

The *demand curve* of an MP for a given reference string is a function  $y = f(x)$  specifying the "swapping load"  $y$  that results from making a "memory space investment"  $x$ . A point  $(x, y)$  of this function is called a *demand point*. The swapping load is represented as a linear form

$$y = m + A\mu,$$

where  $m$  is the *miss rate*, the number of segment faults per unit virtual time, and  $\mu$  is the *information flow rate*, the number of data units per unit virtual time being

Table I. Contributions Associated with Intervals Between References.

	References to segment $i$			Contributions to		
	prior	present	Residence Interval	$\sigma T$	$sT$	$yT$
First	—	$r(t)$	$[t, t]$	$z_i$	$q_i$	$1 + Az_i$
Subsequent	$r(t')$	$r(t)$	$[t' + 1, t'']$	$z_i(t'' - t')$	$q_i + R(i, t'')$	$\begin{cases} 0, & t'' = t \\ 1 + Az_i, & t'' < t \end{cases}$
Post-final	$r(t')$	—	$[t' + 1, t'']$	$z_i(t'' - t' - 1)$	$R(i, t'')$	0

moved to satisfy segment faults. The analyst would normally choose the parameter  $A$  so that  $y$  is proportional to the average time required to service a single swapping operation (queueing for swapping service is excluded). This can be done by setting

$$A = \frac{\text{mean time to transfer one data unit}}{\text{mean access time of secondary store}}$$

Under this interpretation of  $A$ , the total time to complete all the swapping is  $yT$  access times, and the mean swapping delay for one fault of segment  $i$  is  $1 + Az_i$  access times.

There are two possible representations of the memory-space investment  $x$ : Either  $\sigma$ , the *mean resident set size* generated by the MP; or  $s$ , the *mean memory usage cost* (per reference) actually expended by the MP. Notice that  $\sigma T$  is the total (virtual) space-time accumulated among all resident segments; it could be computed by summing resident-set sizes for  $t = 1, \dots, T$ . Likewise,  $sT$  is the total memory-usage cost; it could be computed by adding the total reference cost,  $Q$ , to the total of all retention-cost increments,  $R(i, t) - R(i, t - 1)$ , for all resident  $i$  and  $t = 1, \dots, T$ . However, there are more efficient computational methods than these.

The most efficient methods for measuring the totals  $\sigma T$ ,  $sT$ , and  $yT$  are based on calculating contributions for each *interval* between successive references to a segment [6, 14, 21, 25]. These contributions are summarized in Table I. There are three cases, according as  $r(t) = i$  is a first, an intermediate, or a final reference. A first reference contributes a swap, and memory usage during  $[t, t]$  only. A subsequent reference ends an interval  $[t' + 1, t]$  that spans a pair of successive references; segment  $i$  is resident during a prefix  $[t' + 1, t'']$ , and a swap occurs only if  $t'' < t$ . After a final reference there may be an additional period of residence  $[t' + 1, t'']$ ; in a one-pass measurement, its contribution must be computed after time  $T + 1$  (the procedure cannot discover prior to this time that there are no more references).

In practice, an analyst wishes to evaluate the demand curve of an MP on a given reference string for a given set of threshold values  $\{\Theta_n, n = 1, \dots, N\}$  only. The measurements will yield a corresponding set of demand points  $(x_n, y_n)$ . These points are usually displayed as graphs by connecting adjacent points with straight-line segments. (Fitted interpolation can also be used [26].) This method of display, intended primarily for visual

convenience, has been used for years with page-fault rate functions—e.g., [1, 2, 3, 5, 6, 16, 20, 21, 25]. Mathematically, these graphs approximate a value  $y = f(x)$  by linear interpolation between the nearest pair of measured  $(x, y)$  demand points. If the approximation is too crude, the analyst must calculate demand points for further values of  $\Theta$ .

It is important to remember that the reference and retention costs are abstract quantities used to define memory policies, and that the swapping load does not account for system delays such as queueing for swapping or overhead in placement and replacement. Therefore, costs displayed by demand curves need not correspond to the actual costs of running programs in the system. To assess actual costs, an analyst must use the demand curves to drive simulations or analytic models of a system.

## Generalized Memory Policies

### Generalized Working Sets

The familiar time-window working set for paging,  $W(t, \tau)$ , comprises all pages which have been referenced in the virtual time interval  $(t - \tau, t]$ . (See [9].) If  $r(t - u)$  is the latest reference to page  $i$  prior to time  $t$ , then  $u \geq 1$  and page  $i$  is in the working set whenever  $u - 1 \leq \tau$ .


The parameter  $\tau$  can be regarded as a proportionality constant that relates the value of retaining a page in memory to the cost of retrieving it on a page fault. The working set behaves as if  $\tau$  page-seconds of nonreference are as expensive as one page fault; it removes a page as soon as the cost of retaining it begins to exceed the cost of retrieving it.

The generalized working set MP extends this cost-balancing principle. The cost of retaining segment  $i$  in residence from its prior reference until time  $t$  is  $R(i, t)$ . The cost of swapping (retrieving) it is  $1 + Az_i$ . The threshold  $\Theta$  is the constant of proportionality that relates one unit of swapping to one unit of retention cost. The *generalized working set* (GWS),  $W(t, \Theta)$  for  $t = 1, \dots, T$  and  $\Theta \geq 0$  comprises  $r(t)$  plus all segments for which

$$R(i, t) \leq \Theta(1 + Az_i).$$

This definition implies that a segment can be removed from the GWS at a noninteger time; however, the program is always charged a retention cost of exactly  $\Theta(1 + Az_i)$

Fig. 1. Properties of a demand curve.



for a segment so removed from the GWS. It is easy to see that the GWS satisfies the *inclusion property*

$$W(t, \Theta) \subseteq W(t, \Theta + \epsilon) \quad \epsilon > 0$$

for all  $t$ . This observation shows that the GWS policies are contained in the class of all demand-fetching MP's that have a control parameter  $\Theta \geq 0$  and satisfy the inclusion property. The converse is also true—every demand-fetching MP that has a control parameter  $\Theta \geq 0$  and satisfies the inclusion property is equivalent to some GWS.

To see this, let  $M(t, \Theta)$  denote the resident set of such an MP at time  $t$ , given that its control parameter is fixed at  $\Theta$ . Suppose that the inclusion property holds—i.e.  $M(t, \Theta) \subseteq M(t, \Theta + \epsilon)$ . We will define a retention-cost function  $R(i, t)$  so that the GWS  $W(t, \Theta)$  is identical to  $M(t, \Theta)$ . Since the inclusion property holds, we can imagine that the  $p$  segments are placed on the interval  $[0, \infty]$  so that, for every  $\Theta$ , exactly the segments of  $M(t, \Theta)$  are contained in the interval  $[0, \Theta]$ . At  $t = 0$ , all segments are at infinity, since  $M(0, \Theta)$  is empty. (This is a continuous form of the “stack” [20, 6].) Let  $R(i, t)$  denote the distance of segment  $i$  from the origin; remember that  $R(i, t)$  may depend on some total internal state of (a model of) the program. This distance function satisfies the three properties of retention cost:

- Prior to the first reference to segment  $i$ ,  $R(i, t)$  is infinite, else  $i$  would be in  $M(t, \Theta)$  for some finite  $\Theta$ .
- If  $r(t) \neq i$ , it is impossible for  $R(i, t) < R(i, t - 1)$ . For, if so, segment  $i$  would enter  $M(t, R(i, t))$  at time  $t$ , contradicting the assumption that, for every  $\Theta$ , MP fetches missing segments only when they are referenced.
- If  $r(t) = i$ ,  $R(i, t^+)$  must be 0, else segment  $i$  could not be guaranteed to be in  $M(t, \Theta)$  for every  $\Theta \geq 0$  just after a reference to it.

The foregoing arguments define the sense in which the GWS is a model for the entire class of demand-fetching MP's that have a single control parameter  $\Theta \geq 0$  and satisfy the inclusion property. No MP in this class displays “anomalous behavior,” which would be a de-

crease in  $x$  or an increase in  $y$  when  $\Theta$  is decreased [17].

For each value of  $\Theta$ , the GWS produces values of mean resident-set size  $\sigma(\Theta)$ , mean memory usage cost  $s(\Theta)$ , miss rate  $m(\Theta)$ , and information flow rate  $\mu(\Theta)$ . The inclusion property implies that  $s(\Theta)$  and  $\sigma(\Theta)$  are nondecreasing in  $\Theta$ . Moreover, the segment faults at threshold  $\Theta + \epsilon$  are a subsequence of those at threshold  $\Theta$ ; this implies that  $m(\Theta)$  and  $\mu(\Theta)$  are nonincreasing in  $\Theta$ . Figure 1 summarizes these facts for a demand curve  $y(\Theta) = f(s(\Theta))$ . When  $\Theta = 0$ , it will be true that

$$s(0) \geq Q/T \quad \text{and} \quad y(0) \leq 1 + AV/T,$$

with equality only if there are no repeated references, and only if retention cost is never 0 except for the infinitesimal interval immediately following a reference. When  $\Theta = \infty$ , segment faults occur only at first references; there are  $p$  such faults and they move  $Z = z_1 + \dots + z_p$  data units, whence  $y(\infty) = (p + AZ)/T$ .

## Examples

The *time-window working set* removes a segment which has been unreferenced for  $\Theta$  time units. This effect occurs when  $R(i, t) = (u - 1)(1 + Az_i)$ , where  $r(t - u)$  is the most recent reference to segment  $i$  prior to time  $t$ . For paging, this GWS resembles the original paged working set with  $\tau = \Theta + 1$ .

The *space-time working set* uses  $q_i = z_i$  and sets retention cost to measure the space-time accumulated by a segment after its prior reference. Thus  $R(i, t) = (u - 1)z_i^u$ , where  $r(t - u)$  is the most recent reference to segment  $i$  prior to time  $t$ . Putting all  $z_i = 1$  reduces this GWS to the time-window working set for paging. It is important to note that  $s = \sigma$ —i.e., the mean cost is the same as the mean resident set size when memory usage is measured by space-time.

The function  $R(i, t) = D(i, t) - 1$ , where  $D(i, t)$  is the stack distance of page  $i$  at time  $t$  under a given stack algorithm, is also a valid retention-cost function [20, 6]. Therefore  $W(t, \Theta)$  is precisely the resident set of size  $\Theta + 1$  of the given stack algorithm. In this case  $s(\Theta)$  is the mean stack distance over the referenced pages.

It is possible to specify a program behavior which reflects the GWS [28]. Denote by  $P(i, t)$  the probability of referencing page  $i$  at time  $t$ ; assume that  $r(t) \neq i$  implies  $P(i, t + 1) \leq P(i, t)$  and that  $r(t) = i$  implies  $P(i, t^+) = 1$ . Then  $R(i, t) = 1 - P(i, t)$  is a valid retention-cost function and, for  $0 \leq \Theta \leq 1$ ,  $W(t, \Theta)$  is  $r(t)$  plus all pages whose reference probability is at least  $1 - \Theta$ . A similar idea was suggested in [7].

The page-fault frequency policy (PFF) [5] is not a GWS. The PFF retention costs increase with time, but are reset to zero on any page fault; PFF thus violates retention-cost Property (b). PFF violates the inclusion property and exhibits anomalous behavior [17].

## Generalized Optimal Policies

The purpose of this subsection is to define a demand-fetching MP whose convex demand curve divides the  $(x, y)$  plane into a feasible upper part and an infeasible lower part. This policy will be called the *generalized optimal policy* (GOPT) because no memory policy can generate a demand point below its demand curve. Because the GOPT has lookahead, it is useless for optimal memory management in real time. However, its demand curve, which is easily computed as a byproduct of the GWS's demand curve, can be valuable in showing the analyst how well a program or memory policy behaves.

Recall that memory-space investment  $(x)$  is either the mean size of the resident set  $(\sigma)$  or the mean of memory usage costs  $(s)$ , and that, in the space-time, memory usage cost  $s = \sigma$ . This means that, to find the minimum possible swapping for a given mean resident set size, the analyst needs to study only the demand curve of space-time GOPT. It also means that we may study GOPT only in the  $(s, y)$  plane without loss of generality.

Underlying the GWS is the concept that the threshold  $\Theta$  is the value of one unit of swapping relative to one unit of memory usage. This means that  $s + \Theta y$  can be interpreted as the "net cost" of demand point  $(s, y)$ . The concept underlying GOPT is to make replacement decisions to minimize the "net cost" relative to the given measures of memory usage and swapping.

It now follows that GOPT must remove a segment from the resident set just after its final reference, for any delay would increase memory usage  $(s)$  without affecting swapping  $(y)$ . Indeed, if GOPT opts to remove *any* segment from the resident set, it must do so immediately after a reference to that segment; any delay would increase  $s$  without affecting  $y$ .

It follows from these properties that, for each reference  $r(t) = i$ , GOPT makes just one of two decisions: *Retain*  $i$  until its next reference  $r(t + u)$ , or *remove*  $i$  just after time  $t$ . For a given value of threshold  $\Theta \geq 0$ , the retain decision is taken if and only if

$$R(i, t + u) \leq \Theta(1 + Az_i).$$

Because  $\Theta$  specifies the value of retention relative to swapping, this criterion causes GOPT to select the cheaper decision for each reference. It follows that GOPT minimizes the total cost  $sT + \Theta yT$ .<sup>1,2</sup> Notice that an equivalent

<sup>1</sup> Another way to see this is to consider the effect, on the sum  $s + \Theta y$ , of changing a "remove" to a "retain" decision, and vice versa. Changing the reference  $r(t) = i$  from a "remove" to a "retain" changes the total memory usage cost to  $sT + R(i, t + u)$ , and the swapping cost to  $yT - (1 + Az_i)$ . This changes the total net cost to  $sT + \Theta yT + [R(i, t + u) - \Theta(1 + Az_i)]$ . Since  $r(t)$  is retained under GOPT, the bracketed term is positive—such a change cannot lower the net cost. A similar argument shows that changing GOPT "retain" to a "remove" cannot lower the net cost.

<sup>2</sup> GOPT need not be optimal among nondemand optimal MP's. Let  $F(i, w)$  denote the swapping cost when segment  $i$  is fetched  $w$  time units prior to a reference  $r(i)$ . Note that  $F(i, 0) = \Theta(1 + Az_i)$ . Prefetching would be advantageous if  $R(i, t) - R(i, t - w) + F(i, w) < F(i, 0)$  for some  $w > 0$ .

statement of the GOPT decision rule is: Take the "retain" decision for  $r(t) = i$  only if the *cost/swap ratio*  $R(i, t + u)/(1 + Az_i)$  does not exceed the threshold. If  $r(t)$  is a final reference, then setting  $R(i, t + u)$  to be infinite forces the "remove" decision.

It is easy to see that GOPT satisfies the inclusion property on its resident sets—at threshold,  $\Theta$ , it takes a subset of the retain decisions which it assumes at threshold  $\Theta + \epsilon$ . Therefore, there exists a GWS that simulates GOPT. (One possible GWS uses the retention-cost function  $R'(i, t)$ , defined as follows: Whenever  $r(t)$  and  $r(t + u)$  are successive references to segment  $i$ , set  $R'(i, t') = R(i, t + u)$  for all  $t'$  in the interval  $[t + 1, t + u]$ .)

We can show now that GOPT's demand curve is convex and divides the  $(s, y)$  plane into a feasible upper part and an infeasible lower part. Let  $\rho_0 = 0$ , and let  $\rho_k$  denote the  $k$ th largest of the cost/swap ratios occurring in the reference string. Let  $K \leq T - p$  be the number of distinct finite values of these ratios (the  $p$  final references have infinite ratios). Then  $\rho_0 \leq \rho_1 < \dots < \rho_K$ . When  $\Theta = \rho_k$ , GOPT generates a demand point  $(s_k, y_k)$  for which  $s_k$  is the mean reference cost  $Q/T$  plus the mean of retention costs over all references whose cost/swap ratios do not exceed  $\rho_k$ , and  $y_k$  is the mean of swapping loads over all references whose ratios exceed  $\rho_k$ . As shown in Figure 2, the slope of the line connecting adjacent demand points for  $\Theta = \rho_{k-1}$  and  $\Theta = \rho_k$  is

$$\Delta y / \Delta s = (y_k - y_{k-1}) / (s_k - s_{k-1}) = -1 / \rho_k.$$

This is because all those references which GOPT changes from "remove" decisions to "retain" decisions when  $\Theta$  changes from  $\rho_{k-1}$  to  $\rho_k$  have the same value of cost/swap ratio ( $\rho_k$ ). Note also that for  $\rho_{k-1} \leq \Theta < \rho_k$ , GOPT generates the one demand point  $(s_{k-1}, y_{k-1})$ . Finally, note that no MP can generate a demand point below the line connecting the two demand points for  $\Theta = \rho_{k-1}$  and  $\Theta = \rho_k$ ; for if  $(s, y)$  were such a demand point, its "net cost" would be  $s + \rho_k y < s_k + \rho_k y_k$ , in contradiction to the optimality of GOPT.

Repeating the argument for  $k = 1, 2, \dots, K$  and accounting for the boundary conditions at  $\Theta = 0$  and  $\Theta = \infty$ , we find that the GOPT demand curve is convex and divides the  $(s, y)$  plane into a feasible and infeasible region, as shown in Figure 3.

Suppose that GOPT demand points are computed for a finite set of arbitrary thresholds  $\Theta_1, \dots, \Theta_N$ . The resulting  $N$  demand points will be a subset of the  $K$  possible ones, and the piecewise-linear curve connecting adjacent points will be convex. However, if  $N < K$ , this demand curve may not partition the plane into a feasible and infeasible region.

It is possible to define other optimal MP's based on criteria such as "minimize  $s$  for each given  $y$ " or "minimize  $y$  for each given  $s$ ." Because such MP's may select arbitrary subsets of references to be "retain" and "remove" decisions, they may generate as many as  $2^T$  distinct demand points. However, the demand curve of

Fig. 2. Adjacent demand points of GOPT.





Fig. 3. Infeasible region of demand points.



such an MP must lie on or above the demand curve of GOPT, and need not be convex.

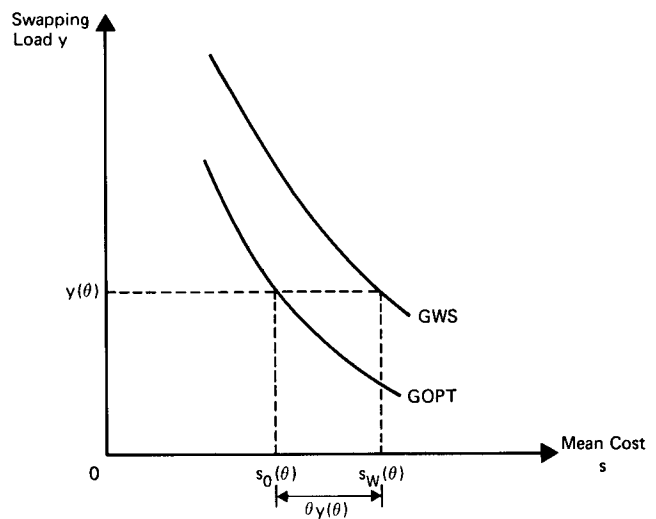
### Examples

Let  $[t, t + u]$  denote an interreference interval of segment  $i$ . When retention cost is measured by space-time,  $R(i, t + u) = (u - 1)z_i$ ; this GOPT retains  $i$  only if  $u \leq 1 + \Theta(A + 1/z_i)$ . For paging with all  $z_i = 1$  and with  $A = 0$ , this GOPT reduces to the vMIN policy [21].

We noted that GWS simulates a stack-paging algorithm if the retention-cost function is the stack algorithm's distance function. However, GOPT is *not* the MIN policy in this case [3, 4]. MIN optimizes  $(\sigma, m)$  demand points over the entire class of fixed-space stack algorithms and, hence, over the entire class of possible stack distance functions; in contrast, GOPT optimizes relative to a single, given stack distance function. Moreover, vMIN may produce a demand curve below that of MIN [21].

An argument similar to the one used to prove the GOPT can also be used to prove that the time-window working set may be optimal among nonlookahead policies, when the program has sufficient locality of reference [10]. The required conditions seem to hold in practice [18].

Fig. 4. Comparison of GOPT and GWS.



### A Relation Between GWS and GOPT

Let  $[t, t + u]$  denote an interreference interval of segment  $i$ . When  $R(i, t + u) \leq \Theta(1 + Az_i)$  both policies retain  $i$  during the interval  $[t, t + u]$ . Otherwise, GOPT removes  $i$  at the beginning of the interval while GWS retains it until its retention cost attains  $\Theta(1 + Az_i)$ . Therefore, for given  $\Theta$ , GWS and GOPT have identical fault sequences; they produce the same swapping load.

The memory-usage cost difference between GWS and GOPT is estimated easily (Figure 4). After each nonfinal reference at which GOPT removes segment  $i$ , GWS generates the additional retention cost  $\Theta(1 + Az_i)$ . After the final reference, GWS may be forced to remove  $i$  at time  $T + 1$ , so that the final cost difference is at least 0 but at most  $\Theta(1 + Az_i)$ . By associating the final GWS cost contribution for segment  $i$  with the initial fault for segment  $i$ , we see that (a) the total of all cost differences cannot exceed the sum of  $\Theta(1 + Az_i)$  contributions at faults—i.e.,  $T\Theta y(\Theta)$ ; and (b) because the cost differences after final references cannot total more than  $\Theta(p + AZ)$  among the  $p$  segments of total volume  $Z$ , the total cost difference is at least  $T\Theta y(\Theta) - \Theta(p + AZ)$ . Thus

$$\Theta y(\Theta) - (p + AZ)/T \leq s_w(\Theta) - s_0(\Theta) \leq \Theta y(\Theta).$$

For small  $\Theta$  or large  $T$ ,  $\Theta y(\Theta)$  is a good approximation to the cost difference.

### Efficient Computation of Demand Curves

Let  $\Theta_0 = 0$  and suppose that  $\Theta_1, \dots, \Theta_N$  is a sequence of increasing threshold values for which GWS and GOPT demand points are to be computed. Often  $\Theta_{n+1} = 2\Theta_n$  for  $n > 0$  gives clear resolution of a demand curve, whence  $N$  is approximately  $\log_2 T$  for a reference string of length  $T$  [26]. In the following,  $[t - u, t]$  will denote

an interference interval of segment  $i$ ; if  $r(t)$  is a first reference,  $u$  will be a large value. The length  $u$  can be computed simply if each segment's time of most recent reference is kept in a table [6, 16, 26].

The four measures ( $m, \mu, s, \sigma$ ) will be specified from information obtained in one pass over the reference string and stored in four sets of  $N + 2$  counters. For  $n = 1, \dots, N$ , the counters are defined as follows:

$a(n)$ , the total swapping load from segment faults that would be saved by increasing  $\Theta$  from  $\Theta_{n-1}$  to  $\Theta_n$ ;

$b(n)$ , the total reference volume for faults tallied in  $a(n)$ ;

$c(n)$ , the total of retention costs that would be added by increasing  $\Theta$  from  $\Theta_{n-1}$  to  $\Theta_n$ ; and

$d(n)$ , the total of resident segment space-times that would be added by increasing  $\Theta$  from  $\Theta_{n-1}$  to  $\Theta_n$ .

For  $n = 0$ , the counters record events for  $\Theta = 0$ . For  $n = N + 1$ , they record all events for  $\Theta > \Theta_N$ . Two additional counters,  $V$  and  $Q$ , tally the total reference volume and the total of (unavoidable) reference costs; each reference to segment  $i$  contributes  $z_i$  to  $V$  and  $q_i$  to  $Q$ . Initially, all the counters contain zeroes.

The values in the counters are updated for  $t = 1, \dots, T$  as follows. If  $R(i, t) = 0$ , set  $n$  to 0, otherwise find the largest  $n$  ( $1 \leq n \leq N + 1$ ) such that  $\Theta_{n-1} < R(i, t)/(1 + Az_i)$ . Then add  $1 + Az_i$  to  $a(n)$ ,  $z_i$  to  $b(n)$ ,  $R(i, t)$  to  $c(n)$ , and  $z_i(u - 1)$  to  $d(n)$ . For initial references,  $R(i, t)$  is chosen to be larger than  $\Theta_N(1 + Az_i)$  for all  $i$ . When  $R(i, t)$  depends only on  $u$ , this can be satisfied by choosing a sufficiently large initial value for  $u$ .

The counter-updating actions will fail to record contributions occurring after the final reference to a segment. This does not affect GOPT, which removes every segment after its final reference; but it does affect GWS. Since the GWS behavior following the final reference to segment  $i$  depends on whether  $R(i, T + 1) \leq \Theta(1 + Az_i)$ , the corrections for these "end effects" are computed by performing the counter updating actions as if  $i = r(T + 1)$  for each segment  $i = 1, \dots, p$ . (See Table I.) We let  $a^*(n)$  denote the total corrections generated for counter  $a(n)$ ; similarly for  $b^*(n)$ ,  $c^*(n)$ , and  $d^*(n)$ .

### Miss Rate, Flow Rate, and Swapping Load

The miss rate, information flow rate, and swapping load are the same for both GWS and GOPT. Reference  $r(t) = i$  produces a fault, whose retrieval demand is  $1 + Az_i$  if and only if  $\Theta < R(i, t)/(1 + Az_i)$ ; it follows that

$$y(\Theta_n) = (a(n + 1) + \dots + a(N + 1))/T.$$

The flow rate is, similarly,

$$\mu(\Theta_n) = (b(n + 1) + \dots + b(N + 1))/T.$$

The miss rate can be calculated as

$$m(\Theta_n) = y(\Theta_n) - A\mu(\Theta_n).$$

### GOPT Mean Cost and Mean Resident Set Size

The GOPT mean cost is denoted by  $s_0(\Theta)$  and its mean resident-set size by  $\sigma_0(\Theta)$ . The total GOPT cost,  $Ts_0(\Theta)$ , is  $Q$  plus all retention costs generated on interference intervals  $[t - u, t]$  for which  $R(i, t)/(1 + Az_i) \leq \Theta$ :

$$s_0(\Theta_n) = (Q + c(0) + \dots + c(n))/T.$$

The total space-time of resident segments,  $T\sigma_0(\Theta)$ , is reference volume  $V$  plus all additional space-time from retained segments:

$$\sigma_0(\Theta_n) = (V + d(0) + \dots + d(n))/T.$$

The large retention cost assumed for an initial reference causes counters  $c(N + 1)$  and  $d(N + 1)$  to receive meaningless values on first references; however, these counters are not used.

### GWS Mean Cost

The GWS mean cost is denoted by  $s_w(\Theta)$ . We noted earlier that  $s_w(\Theta)$  is approximated well by  $s_0(\Theta) + \Theta y(\Theta)$ , whenever  $e = \Theta(p + AZ)/T$  is small compared to  $\Theta y(\Theta)$ . The exact GWS mean cost is calculated by correcting the lower bound on the cost difference,  $\Theta y(\Theta) - e$ , with the additional retention cost contributions following the final references. Let  $s_w^*(\Theta)$  denote the correction; then

$$s_w(\Theta) = s_0(\Theta) + \Theta(y(\Theta) - (p + AZ)/T) + s_w^*(\Theta).$$

At time  $T + 1$ , segment  $i$  contributes  $R(i, T + 1)$  to the correction if  $R(i, T + 1)/(1 + Az_i) \leq \Theta$ ; otherwise it contributes  $\Theta(1 + Az_i)$ . Summing these contributions,

$$s_w^*(\Theta_n) = (c^*(0) + \dots + c^*(n))/T + \Theta_n(a^*(n + 1) + \dots + a^*(N + 1))/T.$$

Since the sum of all the corrections cannot exceed  $\Theta(p + AZ)$ ,

$$s_w^*(\Theta_n) \leq \Theta_n(p + AZ)/T.$$

There is another way of computing the exact  $s_w(\Theta)$ . After the counters have been used to compute the GOPT demand measures, the corrections are added to them. Then the mean cost is computed directly, using corrected  $a$ - and  $c$ -counters, from

$$s_w(\Theta_n) = (Q + c(0) + \dots + c(n))/T + \Theta_n(a(n + 1) + \dots + a(N + 1) - (p + AZ))/T.$$

The quantity  $p + AZ$  is deducted from the corrected  $a$ -counters because there are no retention costs prior to the first references.

Consider a page reference string (all  $z_i = 1$ ) with  $A = 0$  and suppose that the thresholds are chosen to be the

Table II. Example of the Time-Window Working Set

$r(i)$ :	E	C	B	E	A	B	D	C	D	E	B	C	B	B
$R$ :	$\infty$	$\infty$	$\infty$	2	$\infty$	2	$\infty$	5	1	5	4	3	1	0
$b(R)$ :	5	3	2	5	1	2	4	3	4	5	2	3	2	2
$d(R)$ :	—	—	—	10	—	4	—	15	4	25	8	9	2	0

Seg	Size	End Corrections		
		$R$	$b^*$	$d^*$
A	1	9	1	9
B	2	0	2	0
C	3	2	3	6
D	4	5	4	20
E	5	4	5	20

$R$  = Retention-cost value (unreferenced time)  
 $a(\Theta)$  = Number of references of  $R = \Theta$   
 $b(\Theta)$  = Total of segment sizes among references of  $R = \Theta$   
 $c(\Theta)$  = Total of  $R$  values among references of  $R = \Theta$   
 $d(\Theta)$  = Total of  $Rz_i$  among references of  $R = \Theta$   
 $V$  = Total of  $b$ -counters = 43

$R = \Theta$	$a(\Theta)$	Misses $Tm(\Theta)$	$b(\Theta)$	Volume $T\mu(\Theta)$	corrected		GWS $T\sigma_w(\Theta)$	$d(\Theta)$	GOPT $T\sigma_0(\Theta)$
					$d(\Theta)$	$b(\Theta)$			
0	1	13	2	41	0	4	43	0	43
1	2	11	6	35	6	6	82	6	49
2	2	9	7	28	20	10	115	14	63
3	1	8	3	25	9	3	138	9	72
4	1	7	2	23	28	7	158	8	80
5	2	5	8	15	60	12	171	40	120
9	0	5	0	15	9	1	175	0	120
$\infty$	5		15			15			

Table III. Example of the Space-Time Working Set.

$r(i)$ :	E	C	B	E	A	B	D	C	D	E	B	C	B	B	End Corrections		
$R$ :	$\infty$	$\infty$	$\infty$	10	$\infty$	4	$\infty$	15	4	25	8	9	2	0	Seg	Size	$R = c^*$
$Z_i$ :	5	3	2	5	1	2	4	3	4	5	2	3	2	2	A	1	9
															B	2	0
															C	3	6
															D	4	20
															E	5	20

$R = \Theta$	$a(\Theta)$	$c(\Theta)$	Misses $Tm(\Theta)$	$b(\Theta)$	Volume $T\mu(\Theta)$	corrected		GOPT $T\sigma_0(\Theta)$	GWS $T\sigma_w(\Theta)$
						$a(\Theta)$	$c(\Theta)$		
0	1	0	13	2	41	2	0	43	43
2	1	2	12	2	39	1	2	45	67
4	2	8	10	6	33	2	8	53	89
6	0	0	10	0	33	1	6	53	107
8	1	8	9	2	31	1	8	61	123
9	1	9	8	3	28	2	18	70	130
10	1	10	7	5	23	1	10	80	135
15	1	15	6	3	20	1	15	95	155
20	0	0	6	0	20	2	40	95	170
25	1	25	5	5	15	1	25	120	175
$\infty$	5			15		5			

first  $N$  integers (i.e.,  $\Theta_n = n$ ). If retention costs are integers, such as at integer times for the time and space-time retention measures, then  $R(i, t) = n$  implies that one is added to counter  $a(n)$  and  $n$  to counter  $b(n)$  during updating. This implies that  $b(n) = na(n)$ . The mean GWS costs can be expressed as


$$s_w(n) = s_w(n - 1) + m(n)$$

$$+ (a^*(n + 1) + \dots + a^*(N + 1) - p)/T.$$

This generalizes the working-set relation obtained for  $T$  infinite [6, 14]. Since  $m(n)$  and  $a^*(n + 1) + \dots + a^*(N + 1)$  are nonincreasing in  $n$ ,  $s_w(n)$  is increasing and concave downward. A subset of the points on the  $s_w(n)$  curve will define a piecewise-linear curve which is also concave downward.



Fig. 5. Demand curves of example reference string.



### GWS Mean Resident Set Size

In the space-time retention measure,  $s_w(\Theta)$  is the mean resident-set size of GWS. But the mean resident-set size  $\sigma_w(\Theta)$  for an arbitrary retention measure may be more difficult to compute. This is because the additional space-time accumulated, among segments retained by GWS beyond their GOPT removal times, is not related simply to any of our previous measures. To illustrate, let  $v_i$  denote the time required beyond time  $t$  for segment  $i = r(t)$  to accumulate retention cost  $\Theta(1 + Az_i)$ —i.e.,  $R(i, t + v_i) = \Theta(1 + Az_i)$ . Let  $y(i, \Theta)$  denote the swapping load due to segment  $i$ . Ignoring end effects, arguments similar to our previous ones show that

$$\sigma_w(\Theta) \cong \sigma_0(\Theta) + \frac{1}{T} \sum_{i=1}^p v_i z_i y(i, \Theta).$$

This could not be computed unless the  $a$ -counters were partitioned into  $p$  sets, one for each segment.

The exact  $\sigma_w(\Theta)$  can be computed, for the time-window working set, using the available information. For an interference interval  $[t - u, t]$  of segment  $i$ ,  $u > \Theta + 1$  implies that segment  $i$  contributes  $\Theta z_i$  space-time before its removal from the GWS. Reasoning similar to that used before produces

$$\sigma_w(\Theta_n) = (V + d(0) + \dots + d(n))/T + \Theta_n(b(n+1) + \dots + b(N+1) - Z)/T,$$

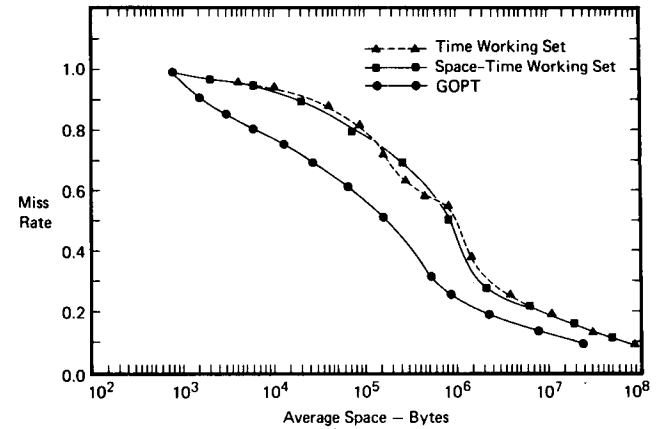
using the corrected  $b$ - and  $d$ -counters. Moreover,

$$\Theta(\mu(\Theta) - Z/T) \leq \sigma_w(\Theta) - \sigma_0(\Theta) \leq \Theta\mu(\Theta).$$

### Examples

Tables II and III show the distributions and calculations of the various measures for a short reference

Fig. 6. Demand curves for database system.



string. The demand curves  $m = f(s)$  are plotted in Figure 5. (Note  $A = 0$  in this case.) For the time-window GWS, we calculated the resident-set sizes to enable a direct comparison with costs in the space-time cost structure. Some of space-time GOPT's demand points are more favorable than for time-window GOPT, since the latter does not necessarily produce the smallest miss rate for a given mean resident set size.

Figure 6 shows the demand curves from an actual segment-reference string obtained from a database system at the IBM San Jose Research Laboratory. The database contains several hundred thousand segments, whose sizes range from tens to hundreds of bytes with an average of about 75 bytes. (See also [22, 23].) The segment-reference string records only references to data segments during several hours of tracing the system; it contains nearly two million references to 183,000 distinct segments, made jointly by several concurrent users. For this system, the time-window and space-time working sets give nearly the same performance, with neither showing a consistent advantage. At high miss rates, they require 15 to 20 times as much space as the optimal policy; this difference reduces to a factor of four or five for low miss rates.

### Conclusion

We have extended the working-set concepts to general cost measures and segment-reference strings. The memory-usage costs include the unavoidable cost of all references to each segment and a nondecreasing cost of retaining each segment while unreferenced. The swapping load is proportional to the delay in retrieving a missing segment. Using threshold  $\Theta$  as the relative value between one unit of swapping and one unit of retention, the generalized working set (GWS) defines the resident set to be the segment referenced at time  $t$  plus all others whose retention-cost to swapping-load ratio does not yet exceed  $\Theta$ . Corresponding to this is a generalized optimal

policy (GOPT) which removes a segment just after it is referenced, if the retention-cost to swapping-load ratio will exceed  $\Theta$  by the time of next reference.

Demand curves for the GWS and GOPT policies can be computed in a single scan of the reference string without simulation. These computations can be done with little space if we are willing to determine demand points for a small number of threshold values. For the space-time GWS and time-window GWS, demand points for adjacent values of  $\Theta$  tend to be very close (or identical) when  $\Theta$  is large [1, 2, 5]. Thus, little resolution is lost in constructing piecewise-linear curves connecting computed demand points.

When all segments are of one size and the cost structure is based on space-time, these results reduce to the familiar ones for paging: GWS becomes the conventional time-window working set [9], GOPT becomes VMIN [21]. Preliminary data showed little practical difference between time-window and space-time GWS performance.

Most of our results do not apply if  $\Theta$  can vary at run time. No policy, including one with  $\Theta$ -variation, can generate a demand point below the GOPT curve: Such variation is of no interest for optimal policies. However, it is possible to vary the GWS threshold so that GWS simulates an optimal policy for part of the time; the resulting demand point may lie below the fixed- $\Theta$  GWS demand curve.

We showed that the cost difference between GWS and GOPT on demand curves  $y = f(s)$  is approximately  $\Theta y(\Theta)$ . For programs whose behavior comprises long phases of referencing over associated locality sets, most of the segment faults occur during transitions between phases [10, 15, 18, 19]. For such programs, the easily computed  $\Theta y(\Theta)$  is a possible measure of the intrinsic differences between a lookahead policy, which can anticipate a transition, and a nonlookahead policy.

*Acknowledgments.* We should like to thank Alan P. Batson, R. Stockton Gaines, G. Scott Graham, Kevin C. Kahn, and Alan J. Smith for comments as the many revisions of this manuscript evolved.

Received March 1976; revised December 1977

#### References

1. Bard, Y. Characterization of program paging in a time sharing environment. *IBM J. Res. and Develop.* 17, 5 (Sept. 1973), 387-393.
2. Bard, Y. Application of the page survival index (PSI) to virtual memory system performance. *IBM J. Res. and Develop.* 19, 3 (May 1975), 212-220.
3. Belady, L.A. A study of replacement algorithms for virtual storage computers. *IBM Syst. J.* 5, 2 (1966), 78-101.
4. Belady, L.A., and Palermo, F.P. On-line measurement of paging behavior by the multilevel MIN algorithm. *IBM J. Res. and Develop.* 18, 1 (Jan. 1974), 2-19.
5. Chu, W.W., and Opderbeck, H. "The page fault frequency paging algorithm. Proc. AFIPS 1972 FJCC, Vol. 41, Pt. 1, AFIPS Press, Montvale, N.J., pp. 597-609.
6. Coffman, E.G. Jr., and Denning, P.J. *Operating Systems Theory*. Prentice-Hall, Englewood Cliffs, N.J., 1973.
7. Coffman, E.G. Jr., and Ryan, T.A. A study of storage partitioning using a mathematical model of locality. *Comm. ACM* 15, 3 (March 1972), 185-190.

8. Courtois, P.J. *Decomposability*. Academic Press, New York, 1977.
9. Denning, P.J. The working set model for program behavior. *Comm. ACM* 11, 5 (May 1968), 323-333.
10. Denning, P.J. Optimal multiprogrammed memory management. In *Current Trends in Programming Methodology III*, K.M. Chandy and R. Yeh, Eds., Prentice-Hall, Englewood Cliffs, N.J., (1978), 298-322.
11. Denning, P.J. The computation and use of optimal paging curves. Tech. Rep. CSD-TR-154, Comptr. Sci. Dept., Purdue University, W. Lafayette, Ind., June 1975.
12. Denning, P.J., and Graham, G.S. Multiprogrammed memory management. *Proc. IEEE* 63, 6 (June 1975), 924-939.
13. Denning, P.J., Kahn, K.C., Leroudier, J., Potier, D., and Suri, R. Optimal multiprogramming. *Acta Informatica* 7, 2 (1976), 197-216.
14. Denning, P.J., and Schwartz, S.C. Properties of the working set model. *Comm. ACM* 15, 3 (March 1972), 191-198.
15. Denning, P.J., and Kahn, K.C. A study of program locality and lifetime functions. Proc. 5th ACM Symp. on Oper. Syst. Principles, Nov. 1975, pp. 207-216.
16. Easton, M. C., and Bennett, B.T. Transient-free working set statistics. *Comm. ACM* 20, 2 (Feb. 1977), 93-99.
17. Franklin, M.A., Graham, G.S., and Gupta, R.K. Anomalies with variable partition paging algorithms. *Comm. ACM* 21, 3 (March 1978), 232-236.
18. Graham, G.S., and Denning, P.J. On the relative controllability of memory policies. Proc. Int. Symp. Comptr. Performance Modeling, Measurement, and Evaluation, Aug. 1977, North-Holland, Amsterdam, 1977, pp. 411-428.
19. Madison, A.W., and Batson, A.P. Characteristics of program localities. *Comm. ACM* 19, 5 (May 1976), 285-294.
20. Mattson, R.L., Gecsei, J., Slutz, D.R., and Traiger, I.L. Evaluation techniques for storage hierarchies. *IBM Syst. J.* 9, 2 (1970), 78-101.
21. Prieve, B.G., and Fabry, R.S. VMIN—an optimal variable space paging algorithm. *Comm. ACM* 19, 5 (May 1976), 295-297.
22. Ragaz, N., and Rodriguez-Rosell, J. Empirical studies of storage management in a data base system. IBM Res. Rep. RJ 1834, IBM Res. Lab., San Jose, Calif., May 1976.
23. Rodriguez-Rosell, J. "Empirical data reference behavior in data base systems. *IEEE Computer* 9, 11 (Nov. 1976), 9-13.
24. Slutz, D.R. A relation between working set and optimal algorithms for segment reference strings. IBM Res. Rep. RJ 1623, IBM Res. Lab., San Jose, Calif., July 1975.
25. Slutz, D.R., and Traiger, I.L. A note on the calculation of average working set size. *Comm. ACM* 17, 10 (Oct. 1974), 563-565.
26. Smith, A.J. A modified working set paging algorithm. *IEEE Trans. Comptrs. C-25*, 9 (Sept. 1976), 907-914.
27. Sreenivasan, K., and Kleinman, A.J. On the construction of a representative synthetic workload. *Comm. ACM* 17, 3 (March 1974), 127-132.
28. Tran-Quoc-Te. An open formulation of working set policies. Rep. 10, Proj. MIMOSA, Facultes Universitaires N.-D. de la Paix, Belgium, Dec. 1975.