Optimal Arrangements of Cartridges in Carousel Type Mass Storage Systems

SRIDHAR SESHADRI*, DORON ROTEM† AND ARIE SEGEV‡

*Department of Statistics and Operations Research, Leonard N. Stern School of Business, New York University, New York, NY 10012

†Information & Computing Sciences Division, Lawrence Berkeley Laboratory, Berkeley, CA 94720 and San Jose State University, Department of MIS, San Jose, CA, USA and

‡Haas School of Business Administration, University of California, and Information & Computing Sciences Division, Lawrence Berkeley Laboratory, Berkeley, CA 94720, USA

Optimal arrangements of cartridges and file partitioning schemes are examined in carousel type mass storage systems using Markov decision theory. It is shown that the Organ-Pipe Arrangement is optimal under different storage configurations for both the anticipatory as well as the non-anticipatory versions of the problem. When requests arrive as per an arbitrary renewal process this arrangement is also shown to minimize the mean queueing delay and the time spent in the system by the requests.

Received September 18, 1993, revised December 11, 1994

1. INTRODUCTION

A large number of applications such as multimedia databases, document retrieval and scientific databases, require high-capacity mass storage systems that can hold multigigabytes or even terabytes of data. These systems contain a hierarchy of storage devices typically consisting of large robotic tape libraries, optical disks and magnetic disks. In such systems, files are automatically migrated across the storage hierarchy such that the more frequently used files reside on the faster and more expensive (in terms of dollars per byte) storage media. In this paper we analyse data placement strategies on a carousel type robotic tape library. Such devices are quite commonly found in mass storage systems, for example Magnus Jukebox Library and Lago Systems LS/300L are carousel type devices that use 8 mm tape cartridges with a total capacity of 270 GB (Ranade, 1992).

The carousel type mass storage system is a configuration found in systems catering for low to medium tertiary storage requirements. The system usually has a number of storage locations for cartridges arranged on the inner periphery of a carousel (see Figure 1). The system responds to a request for loading a cartridge by the movement of the carousel to align the required cartridge in front of a read/write head, and a robot does the actual loading or unloading. The problem addressed in this paper is the optimal allocation of cartridges to the storage locations and files to the cartridges.

A similar problem has been solved recently in the context of determining warehouse storage by Fujimoto (1991). Using a Markovian model, Fujimoto proves the optimality of the Organ-Pipe Arrangement when only one cartridge is stored per location. An Organ-Pipe Arrangement (OPA) is one in which cartridges, or more generally items, are first sorted in the descending order of the probability that they will be requested. The first item is allotted to the central location (for a circular storage device such as a carousel, the choice of this location is arbitrary). Then the remaining items are placed alternatively to the left and right of the central location. The picture made by the graph of the probabilities with respect to locations resembles an organ-pipe. The model solved by Fujimoto is called non-anticipatory based on a terminology introduced in King (1990). In the non-anticipatory case the storage device is not permitted to be repositioned between requests even if time is available for doing so; whereas in the anticipatory case the controller can reposition the device between requests. In Fujimoto's terminology the mass storage carousel is a single dimensional bi-directional system. Singe refers to the number of cartridges per location and the carousel is called bi-directional as it can be rotated in either direction. Fujimoto gives an extensive literature survey in the area of warehousing and also refers to papers on optimal spatial permutations (Bergmans, 1972; Groosman and Silverman, 1973; Yue and Wong, 1973 and Karp *et al.*, 1975). None of these directly address the specific problem of allocation of storage space in carousels. Fujimoto concludes that the closest paper that solves the carousel problem is that of Lim *et al.* (1985), where the optimality of OPA is proved based on Bergmans' analysis. Fujimoto adds a missing step in Bergmans' proof for the case when only one cartridge (item) can be stored per location, and conjectures the optimality of the OPA arrangement for the case when more than one cartridge can be stored per storage location but does not prove the optimality of this policy.

The case of a mass storage system is a bit different, because depending on the load, it is possible to reposition the carousel before another request arrives. This is the

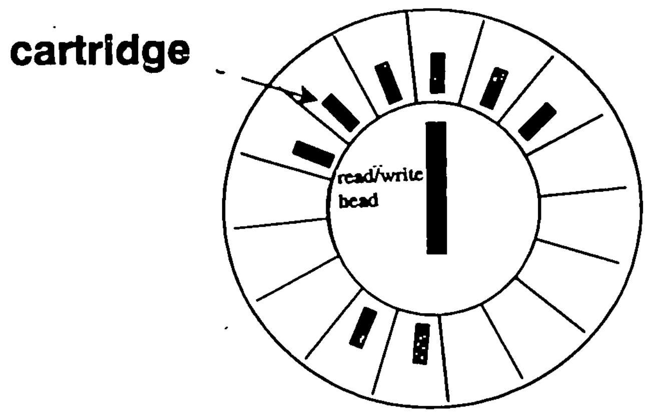

FIGURE 1. A carousel with one read/write head.

anticipatory case as defined above. Anticipatory policies for disk arm control are offered as a good strategy by King (1990). A well known example of such a policy is the greedy policy of 'nearest head' shown to be optimal for a disk with two read heads by Hofri (1983). In that work the head which does not serve the current request is allowed to 'jockey' to an optimal location in anticipation of the next request.

In this paper, we use a Markovian model for the carousel problem, and show that for both the anticipatory and non-anticipatory versions of the problem, the optimal arrangement is the *Organ-Pipe Arrangement* when there is a single read/write head. The results are then extended to the case when two heads are provided. Finally we extend the results to the case where requests arrive as per an arbitrary renewal process and show that for the non-anticipatory case the OPA arrangement minimizes both the mean queueing delay as well as the average time spent in the system by the requests.

2. MODEL

The carousel is modeled as having n storage locations. One cartridge can be placed in each of the locations. We will assume that there are at most n cartridges to be loaded (else the problem will need to incorporate caching). When extending the results to file allocation we will model the cartridge as capable of accommodating at most m files. This is a simplification because file sizes need not be the same and file migration policies could dictate the size distribution of files found in the mass storage system. However the simplification enables the solution of a problem which otherwise will have to contend with random and non-stationary sizes of files and bin packing type of limitations. From experience however, it is seen that storing *larger* files in the tertiary system leads to a better retrieval performance. So the file migration policies may tend to even out the size distribution. In any case, the operating policies described below can be extended to the case of continuous variables representing the items stored and thus are not necessarily of limited usefulness in practice.

In order to model the demand process, it is assumed that the probability a particular cartridge will be requested next is independent of the past requests. Thus the request pattern forms a Markov chain with state probabilities, p_i, the probability that the next cartridge requested will be i and $\Sigma_1^n p_i = 1$. If the number of cartridges is less than n then the remaining probabilities are set to zero. Unless specified, in all cases the number of read/write heads is one. A similar model will be used for the file allocation problem. The Markovian assumption is reasonable considering that the mass storage system does not actively participate in user processing and is more like a library or repository in its functions. In the non-anticipatory case, called case I, the carousel cannot be moved to a specific location in anticipation of a request—whereas in the anticipatory case, termed case II, the carousel can be repositioned between requests. It will be assumed that the travel time is linear and the shortest distance to travel will be realized. The objective throughout is to minimize the mean delay to service a request and except in proposition 7, it is always assumed that requests do not interfere with one another. The last assumption is reasonable for mass storage systems and the ability to reposition in between requests is physically possible but not found in current systems.

Number the locations on the carousel as 1 through n. Assume that cartridge $\pi(i)$ is stored in location i. In case I, as we have a finite state space for the Markov chain, it follows that there exists a unique stationary distribution of the position at which arriving requests find the carousel (see Wolff, 1990 for example). Direct verification shows that the (stationary) probability that a request will find the read head at location i is equal to the probability, $p_{\pi(i)}$, that the cartridge stored at location i will be requested.

In dealing, with the file allocation model in case II, there are two levels of decision. First the files must be allocated to cartridges. Then the cartridges must be arranged in storage locations. Once the file allocation has been carried out, the request probability for a particular cartridge is fixed by the sum of the probabilities of requests for files allotted to the cartridge. It follows that the above stationary distribution holds good once the request probability for cartridges has been computed.

In case I, using the above notation, cartridge $\pi(i)$ is stored in location i, the expected travel distance, $ED(\pi)$, per request is given by (see Wolff, 1990 for example):

$$ED(\pi) = \sum_i p_{\pi(i)} \sum_j p_{\pi(j)} d(i,j)$$

where $d(i,j)$ is the shortest rotational distance between locations i and j. By substituting any function f of the distances $d(i,j)$ in the above formula we also obtain the expected value of that function, i.e.

$$Ef[D_{(\pi)}] = \sum_i p_{\pi(i)} \sum_j p_{\pi(j)} f[d(i,j)]$$

This fact will be used in proposition 1.

In case II, we do not need to use the Markov chain at all. This is because in the anticipatory case it is assumed that the read head can be positioned very quickly to a given position before a request arrives. Given that requests are independent of one another, the optimal repositioning strategy will be stationary and deterministic. Therefore the read head will always be positioned at the same place before a request arrives; a fact that allows us to search within this class of policies in determining the optimal allocation scheme. Thus, if the read head is always repositioned between requests at location i, then the expected distance traveled per request (ignoring the repositioning distance) will be given by:

$$ED(\pi) = \sum_j p_{\pi(j)} d(i,j)$$

Note that the expected repositioning distance is also equal to the above value.

3. THE OPTIMALITY OF THE ORGAN-PIPE ARRANGEMENT

An Organ-Pipe Arrangement is one in which the cartridges are placed in an alternating arrangement. The cartridges are ranked in descending order as per their request probability. Let cartridge #i be the one with the i^{th} largest request probability and let the storage locations be numbered in clockwise fashion. Then cartridge #1 is placed in location l cartridge #2 in location 2, 3 in location n, #4 in location 3, #5 in location $(n-l)$ etc. This arrangement proves to be optimal under a variety of modeling assumptions as described in this section.

The basic condition for optimality of an arrangement is obtained by imagining that a line of symmetry (of any orientation) is drawn across the carousel, see Figure 2. Let the sum of request probabilities for cartridges on the left side of this line be larger than the sum on the right [ignoring the location(s) bisected]. Let i and j be the cartridges stored in mirror image location on the left and right sides of the line respectively. Then intuitively speaking we expect that the request probability for cartridge i should be larger than for cartridge j in the optimal arrangement. This property is called the pairwise Majorization Property (PMP) by Fujimoto. In case I, in fact a necessary and sufficient condition for an arrangement to be optimal is that PMP holds over all symmetry lines. And not very coincidentally the OPA possesses this property. A sketch of the proof (partly provided by Fujimoto) follows but with a strengthening of the result. The strengthening is in the sense that if we start out initially with the stationary distribution of the Markov chain, then OPA minimizes the distance traveled at each transition in the sense of stochastic order. By definition, if X and Y are random variables, then X is larger or equal to Y in the stochastic ordering sense, denoted by $X \geqslant_{st} Y$, if Prob $(X > t) \leqslant \text{Prob}(Y > t)$ for all t. In Wolff (1990) it is shown that $X \geqslant_{st} Y$ is equivalent to the condition that for any non-decreasing function, f, $E[f(X)] \geqslant E[f(Y)]$.

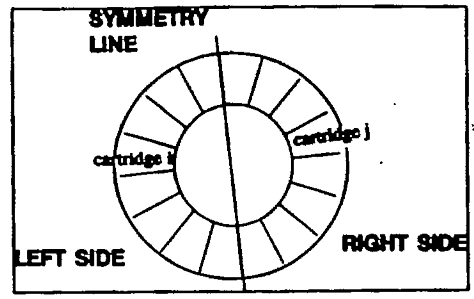

FIGURE 2. The pairwise majorization property.

Proposition 1. (Bergmans, 1972 and Fujimoto, 1991. The OPA arrangement is optimal in case I and also in the stochastic sense described above.

Proof. Let there be a line of symmetry over which the pairwise majorization property is violated. We will call the side that has the higher sum of probabilities as the left side. Assume that the symmetrical dividing line is vertical. Denote a violating assignment to be one where a cartridge on the left side has a lower request probability than the one in the mirror image location on the right. For example, the fourth cartridge on the left of the symmetry line has a smaller request probability compared to the fourth one on the right side of the line.

We will show that this arrangement can be improved by interchanging all pairs of violating assignments simultaneously. This helps because (i) the cost with respect to the un-interchanged cartridges considered by themselves is unchanged, (ii) the cost with respect to the interchanged cartridges by themselves is unaffected, while (iii) the interaction between the interchanged and un-interchanged cartridges leads to lower cost. The case (iii) can be proved by the following argument: fix an un-interchanged cartridge location. The distance by which a cartridge has moved away from this location due to the interchange is exactly the distance by which another cartridge has moved closer to this location. The argument can be formalized by denoting the locations on the left of the symmetry line from which the cartridges must be interchanged to be $\{l1, l2, l3, \ldots, lk\}$ and the mirror image locations on the right (where they must be placed) to be $\{r1, r2, r3, \ldots, rk\}$. Let lu and ru stand for the mirror image locations of two uninterchanged cartridges, on the left and right sides of the symmetry line. Consider the impact of cost in (iii) with respect to lu and ru. The change in cost with respect to these two locations is given by:

$$p_{\pi(lu)}\left[\sum_{j=1}^{k} p_{\pi(l_j)} d(lu, l_j) + \sum_{j=1}^{k} p_{\pi(r_j)} d(lu, r_j)\right] -$$

$$p_{\pi(lu)}\left[\sum_{j=1}^{k} p_{\pi(r_j)} d(lu, l_j) + \sum_{j=1}^{k} p_{\pi(l_j)} d(lu, r_j)\right] +$$

$$p_{\pi(ru)}\left[\sum_{j=1}^{k} p_{\pi(l_j)} d(lu, l_j) + \sum_{j=1}^{k} p_{\pi(r_j)} d(lu, r_j)\right] -$$

$$p_{\pi(ru)}\left[\sum_{j=1}^{k} p_{\pi(r_j)} d(lu, l_j) + \sum_{j=1}^{k} p_{\pi(l_j)} d(lu, r_j)\right]$$
$$= (p_{\pi(lu)} - p_{\pi(ru)}) \sum_{j=1}^{k} (p_{\pi(l_j)} - p_{\pi(r_j)})(d(lu, l_j)$$
$$- d(lu, r_j)) \geqslant 0$$

The last step follows from the fact that $p_{\pi(l_j)} < p_{\pi(r_j)}$ as PMP is violated at the symmetric locations l_j and r_j by assumption, $p_{\pi(lu)} \geqslant p_{\pi(ru)}$ as the cartridges in these two positions satisfied PMP, and $d(lu, l_j) \leqslant d(lu, r_j)$.

Fujimoto argued as above. However the above argument holds when the distances are substituted by any increasing function of the distance. To verify this substitute each distance $d(x, y)$ by the function $f[d(x, y)]$ in the proof given above. This shows that the interchange reduces the distance traveled at each transition in the sense of stochastic order. In our case, this implies that all the arguments carry over to the carousel rotational time rather than rotational distance, because the time is an increasing function of the distance. The importance of starting out with the stationary distribution for demonstrating this ordering must be noted. This artifice does not affect costs averaged over a long period of time, and will be used in a queueing context in proposition 7.

Next it is necessary to show that only an OPA has the PMP over all symmetry lines. Here we deviate from Fujimoto and use a proof by induction. Order the cartridges in decreasing value of request probabilities. Place cartridge #1 in location 1. If cartridge #2 is not placed in location 2 or n then pass a symmetry line adjacent to the location in which #2 has been placed such that #1 is on the opposite side of #2. Let #1 be on the left side. Then that side has to have the larger sum of probabilities else PMP will be violated. But #2 can be interchanged with the cartridge immediately on the left of the symmetry line which is a contradiction. Let the first k cartridges be placed in OPA, starting with #2 in location 2. If the cartridge $\#(k+1)$ is not placed in OPA label the cartridge placed in the OPA location instead as $\#j$. Draw a symmetry line such that $\#j$ and $\#(k+1)$ are in mirror image locations on this line. There are three cases to consider. (i) If $\#j$ is closer to cartridge #1, the sum of probabilities on $\#j$s side of the line is greater. This leads to a contradiction. (ii) If both $\#j$ and $(\#k+1)$ are equidistant from cartridge #1 then k must be an even number. Draw a symmetry line through cartridge #1. The sum of probabilities on $\#j$s side must be higher as k is an even number and so cartridge #2 is on its side of the line. (iii) The case where $\#(k+1)$ is closer to cartridge #1 cannot occur. This completes the proof of the proposition

QED

Proposition 2. The OPA arrangement is optimal in the anticipatory case too.

Proof. In the anticipatory case, by the independence property we will always position the head at the same location between requests (ignoring ties). So each cartridge is essentially assigned a travel distance on a permanent basis. By the inequality of Hardy *et al.* (1991), given two sequences $w_l, \ldots, w_n$ and $p_l, \ldots, p_n$, of all permutations π of the p_i's, the one that minimizes the sum: $\Sigma_{i=1}^{n} w_i p_{\pi(i)}$ is the one that forms each term in the above sum by multiplying the j^{th} largest value of the p_i's with the j^{th} smallest value of w_i's. If we fix the position of the head at a particular location, exactly one cartridge will be at distance 0 (distance measured in units of $2\pi/n$), at most two at distance 1 etc. If we consider our sequence of possible rotational distances as the weights w_i, it follows that the OPA arrangement is optimal and the anticipatory position of the head should be at location i with the maximum $p_{\pi(i)}$.

QED

Remark. The proof assumes that either the move to reposition the head is completed before the next request arrives, or the move is always completed regardless of whether a request has arrived. The latter case assumes no interruptions are allowed during the rotation of the carousel towards its anticipatory position.

In the next proposition we consider the case that each cartridge can hold m equi-sized files and request probabilities are given for each file. This is analogous to Theorem 1.3.1 of Wong83 which deals with records and pages, however the method of proof here is different as we use a direct interchange argument whereas Wong83 uses Schur functions.

Proposition 3. The OPA policy is optimal for the non-anticipatory case and when there are m files stored per cartridge.

Proof. The OPA arrangement in this case is obtained by sorting mn files in descending order of request probabilities and grouping the first m in bundle #1, the second m in bundle #2 etc. Then the bundles are placed in OPA fashion. The proof of optimality is a straightforward extension of proposition 1. Let the OPA property not hold for the optimal arrangement. But by proposition 1 the property has to hold with respect to the *bundles* of files. Without loss of generality let the property *not* hold for *files* in two mirror image locations say across a given symmetry line. Also let the sum of probabilities on the left side of the symmetry line (as before) be larger. Call these mirror image locations on the left side i and that on the right j. What the violation implies is that the probability that a cartridge in location i will be requested can be made larger by interchanging files between locations i and j. Without loss of generality, label the files in location i as l through m and those in location j as $(m+l)$ through $2m$. Label the sorted order of these $2m$ files in descending value of request probability as l' through $2m'$. Consider bundling l' through m' into location i and the rest in location j. Note that this interchange does not affect the cost of *interaction* between other locations. Also the interactions between other locations and interchanged files leads to a decrease in cost as in (iii) of proposition 1. We thus are left with comparing

$$\sum_{i=1}^{m} p_i \sum_{i=m+1}^{2m} p_i \quad \text{and} \quad \sum_{i=1}^{m'} p_i \sum_{i=m'+1}^{2m'} p_i.$$

The minimality of the second expression obtained from a simple interchange argument. To see this, let $p_l < p_{m+1}$. Then interchanging these two alone leads to the difference:

$$-p_{m+1}\sum_{i=2}^{m} p_i - p_l \sum_{i=m+2}^{2m} p_i + p_{m+1} \sum_{i=m+2}^{2m} p_i + p_l \sum_{i=2}^{m} p_i$$

$$= (p_{m+1} - p_l)\left(\sum_{i=m+2}^{2m} p_i - \sum_{i=2}^{m} p_i\right) \tag{1}$$

$$\sum_{i=1}^{m} p_i - \sum_{i=m+1}^{2m} p_i \geqslant 0 \quad \text{and}$$

$$p_{m+1} > p_l \Rightarrow \sum_{i=m+2}^{2m} p_i - \sum_{i=2}^{m} p_i < 0 \tag{2}$$

(1) and (2) together with the hypothesis shows that the interchange of the two files is beneficial. The proof then is completed by showing that only by sorting the files and bundling them can we avoid any violation of the OPA arrangement as defined initially. But this is easy, because (i) the files are first sorted and so the bundles have descending sums of probabilities. So if two mirror image locations violate the OPA property with respect to files then *all* the files in the two locations need to be interchanged. But the bundles are in OPA order leading to a contradiction. And (ii) if the files were not sorted then an interchange is always possible.

QED

Proposition 4. The OPA order is optimal for the anticipatory case when m files can be placed per cartridge.

Proof. Similar to proposition 2.

Proposition 5. When there are two read heads placed symmetrically opposite one another, the optimal anticipatory policy is OPA on *each half* of the carousel.

Proof. We need to define this ordering and will do so shortly. The basic idea is that the repositioning is the same between requests. So the anticipatory action allots *a permanent distance to be traveled* to each file. If the

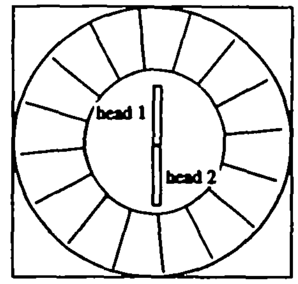

FIGURE 3. A carousel with two read/write heads symmetrically positioned.

number of storage locations is even, say $2x$, then the available values for the distance are $2m$ zeros, $2m$ ones,... By the Hardy, Littlewood, Polya inequality an optimal arrangement is to order the files per descending request probability and place files #1 through #m in location 1, files #$(m+1)$ through #$(2m)$ in location $(x+l)$ etc. In between requests, the strategy is to bring the carousel with locations 1 and $(x+1)$ aligned with the heads. The case of odd number of storage locations is similarly solved.

QED

The case of two read heads and non-anticipatory type of operations is more difficult. The problem is which way should the carousel rotate? First, let us assume that there are an even number of cartridge locations on the carousel; in which case the direction of rotation is immaterial as far as the next request is concerned. In the next proposition we show that for even number of cartridge locations the nearest head policy combined with OPA on each half of the carousel will be optimal. Interestingly, this is not optimal for odd number of locations as we show in the next example.

Example. Consider the carousel schema of Figure 4 with five cartridge locations. In this case, to ensure that the heads are aligned with some cartridge locations, the smaller angle between the two heads is $4\pi/5$ (and the big one $6\pi/5$). Let us assume that each of the request probabilities of cartridges at locations 1, 2 and 4 are a and that of cartridges at locations 3 and 5 are each $0.5–3a/2$. The carousel is currently at the position shown in Figure 4a when a request to read cartridge at location 5 arrives. When the value of a tends to zero, it is clear that rotating the carousel to the position of Figure 4c (the black head serving the request) is optimal as the two heads are now positioned such that with probability $1–3a$, the expected rotational distance for serving the stream of future requests is 0. On the other hand, the nearest head policy will lead to oscillations between positions of 4a and 4b at each future step with probability arbitrarily close to 0.5.

Proposition 6. When there are an even number of locations, OPA is optimal.

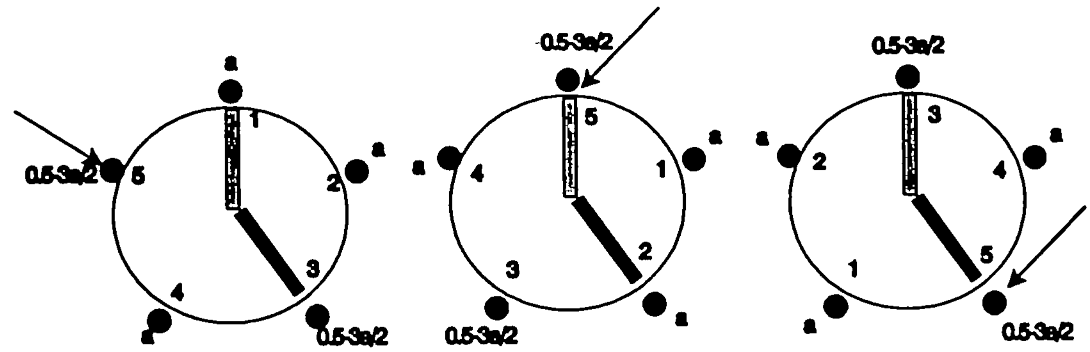

(a) (b) (c)

FIGURE 4. (a) Current position of carousel, when cartridge at location 5 is requested. (b) Nearest head policy, carousel rotates $2\pi/5$ in the clockwise direction, grey head serves the request (c) optimal policy carousel rotates $4\pi/5$ in the counter-clockwise direction, black head serves the request.

Proof. The nearest head policy is the optimal rotational strategy. Therefore the problem can be reduced to the case when there are $n/2$ locations to be filled each capable of holding two cartridges. Proposition 3 provides the result.

QED

Remark. The optimality in propositions 3 and 6 can be extended to hold in the stochastic sense as done in proposition 1.

In many practical multiuser applications, the requests to the mass storage system are queued and serviced on a first come first served basis (FCFS). In the next proposition we utilize two powerful theorems from Stoyan (1983) to show that OPA is optimal in this case as well.

Proposition 7. For the non-anticipatory case, for any of the three models of proposition 1, 3, and 6, when requests arrive as per a general renewal process and they are attended on a first come first served (FCFS) basis, OPA minimizes the average queueing delay as well as the time spent in the system.

Proof. The proof follows from the fact that by propositions 1, 3, and 6, when starting out with the stationary distribution, the time to serve a request is the smallest in the stochastic sense under the OPA order. The rest of the proof can be found in theorems 5.2.1 and 6.2.1 (Stoyan, 1983).

QED

4. SIMULATION RESULTS

In this section we use simulation to (i) investigate the case when successive requests for cassettes can be correlated and (ii) determine how a repositioning (anticipatory) policy performs when there is interference between requests. As a by product of the analysis, we also obtain answers to the questions of how much improvement can be obtained by using the Organ Pipe Arrangement over a random placement of cassettes in the carousel and what is the extent of savings due to the use of an anticipatory policy.

4.1. A Markov chain model for requests probabilities

Successive requests for cassettes will not be independent in any practical situations. A plausible modeling approach in such situations is to use a Markov chain to model the dependence between successive requests. We assume that given the current cassettes requested is i, the probability that the next cassette requested will be j will be given by p_{ij}, where $\Sigma_{j=1}^{n} p_{ij} = 1$. We assume that the transition matrix $\boldsymbol{P} = (p_{ij})$ is irreducible and that the steady state probability vector, $\boldsymbol{p} = (p_1, p_2 \ldots, p_n)$ for the Markov chain is given by solving the equation $\boldsymbol{p} = \boldsymbol{pP}$. It may be verified that if anticipation is not permitted, then the optimal arrangement of cassettes is still OPA based on $\boldsymbol{p}$, and therefore all the previous results for the non-anticipatory case will carry over. But the anticipatory case is very different now, as it is possible to do state dependent anticipation, i.e. depending on what the last request was we can reposition the reading head optimally for the next request. Moreover given the arrangement (of cassettes) there is an optimal anticipation point which can be determined simply by computing the expected cost of travel using the transition probabilities. For example if the cassette i were placed in the position π_i, $i = 1, 2, \ldots, n$, and if the last request was for cassette j, then the optimal repositioning should be done at

$$\min_k \sum_{i=1}^{n} p_{ji} d(\pi_k, \pi_i)$$

where $d(\pi_k, \pi_i)$ is the distance between the locations at which cassette k is placed and the location where cassette i is placed. Unfortunately, determining the optimal arrangement of cassettes is a very hard problem (it can be shown to be in the class of NP-Complete problems). Therefore we tackled the optimal arrangement problem through simulation.

We generated random transition matrices of size *nxn*, $n = 5, 6, 7, 8, 9, 10$. For each matrix of size $n \times n$, we generated n^2 uniformly generated random variables, assigned these to be an $n \times n$ matrix and normalized their sum across rows to be unity. To determine the optimal arrangement of cassettes, we used brute force search over the $(n-l)!$ possibilities. This limited the size of the matrices. To obtain the OPA arrangement for the case when no anticipation was permitted, we solved for the steady state probabilities through recursive convolution of the transition probability matrix. This gave the steady state vector $\boldsymbol{p}$. We also wished to obtain a good lower bound for the minimum cost as well as heuristics that can be applied to larger problems (a typical carrousel would have $n = 40$ cassettes). A simple lower bound can be obtained by assuming that the cassettes can be rearranged between requests and put into the OPA order depending on the state. For example, if the last request was for cassette i, then we use the p_{ij}, $j = 1, 2, \ldots, n$ to rearrange the cassettes. This clearly provides us with a lower bound. We show this lower bound in Tables 1(a–d). We also tested three heuristics: (i) assume that the cassettes are placed in OPA order based on $\boldsymbol{p}$ and reposition the reading head between requests, called OPADYN for dynamic OPA (as against using OPA and a non-anticipatory policy), (ii) use each of the vectors $(p_{ij},\ j = 1, 2, \ldots, n)$, $i = 1, 2, \ldots, n$ to determine n different organ pipe arrangements, compute the expected cost with repositioning between requests being allowed and choose the best of these n arrangements, denoted as GREEDY and (iii) using the OPA based on $\boldsymbol{p}$ compute which request leads to the highest cost, i.e. if the organ pipe arrangement is given by π_i, $i = 1, 2, \ldots, n$, then find

$$\max_j \left\{ p_j \left[\min_k \sum_{i=1}^{n} p_{ji} d(\pi_k, \pi_i) \right] \right\}$$

call this row $\boldsymbol{h}$. Then use an anticipatory policy and OPA based on the transition probabilities, $p_{hj}, j = 1, 2, \ldots, n$. This heuristic is called HICOST. We investigate the case where there is only one read head.

Tables 1a–d summarize our simulations. We conducted 20 trials for $n = 5, 6$ and only 10 trials for the rest. It is seen that the lower bound gets progressively worse as n increases. The GREEDY heuristic performed well in these experiments. In Table 1e we summarize these findings. We see that the GREEDY heuristic performed the best, and that the degree of sub-optimality is on the average <3.2%. The worst case performance was also below 5.5%. The other interesting finding from these experiments was that an anticipatory policy can save on the average 13–17% of travel time. These findings are also summarized in Figure 5 showing the average sub-optimality of the heuristics and the savings from using the best heuristic arrangement and an anticipatory policy over OPA without anticipation.

4.2. Effect of interference of requests on anticipatory policies

In proposition 7, we restricted our attention to non-anticipatory policies. In practical situations it is of interest to know when to use anticipatory policies when requests arrive as per some random process. Repositioning the reading head will add some overhead because of the travel time for moving the head to an anticipatory position, but repositioning saves on subsequent travel time—so when does the trade-off between these two effects favor repositioning? In the commercial systems available today, it is not possible to change the command for moving to a position while the command is being executed. This creates a situation where we cannot change our mind once having given a command to reposition and request arrives before that command is executed. There are several options available for mitigating this effect. One method would be to reposition one location at a time, for example if we are currently at a position i and would like to reposition to location $i+2$, we could do the repositioning in two steps, first to

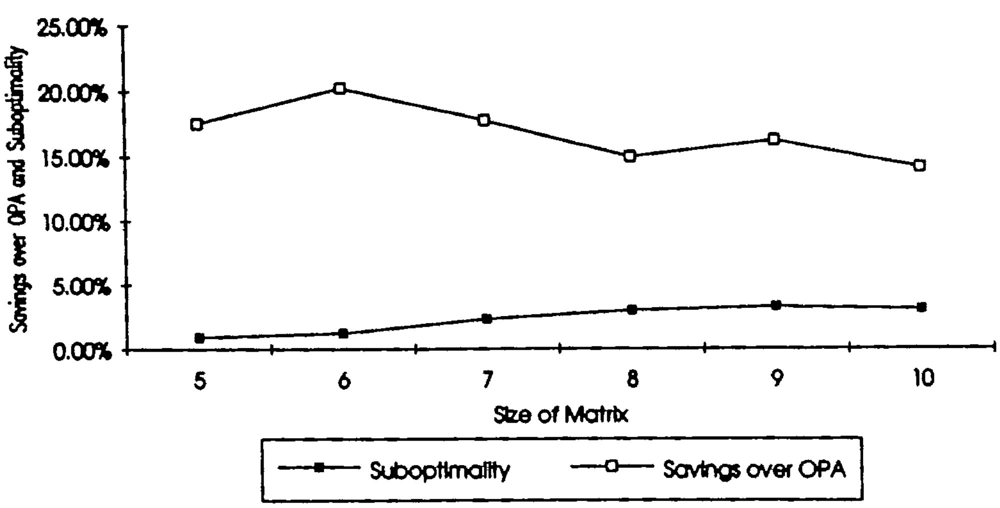

FIGURE 5. Performance of OPAbased heuristic.

TABLE 1(a). 20 Trials for $n = 5$

Trial	OPASTAT	LOWBD	OPADYN	GREEDY	OPTIMAL	HICOST	Min of 3	OPADYN?	GREEDY	HICOST	DEVN %	SAVEOPA
1	0.972455	0.83891	0.918705	0.891628	0.868988	0.891628	0.891628	0	1	1	2.54%	9.07%
2	1.11404	0.828802	0.94817	0.846393	0.846393	0.923408	0.846393	0	1	0	0.00%	31.62%
3	1.06171	0.839621	0.92346	0.881033	0.874855	0.942856	0.881033	0	1	0	0.70%	20.51%
4	0.877076	0.689953	0.743656	0.733139	0.733139	0.734606	0.733139	0	1	0	0.00%	19.63%
5	0.989339	0.89725	0.961437	0.945652	0.92717	0.961437	0.945652	0	1	0	1.95%	4.62%
6	1.00345	0.819481	0.8955	0.851882	0.851882	0.888747	0.851882	0	1	0	0.00%	17.79%
7	1.07319	0.925613	0.966005	0.928981	0.928981	1.00002	0.928981	0	1	0	0.00%	15.52%
8	1.01771	0.880759	0.973741	0.901222	0.901222	0.924178	0.901222	0	1	0	0.00%	12.93%
9	1.08301	0.872763	0.93368	0.906368	0.906368	0.906368	0.906368	0	1	1	0.00%	19.49%
10	0.856971	0.637321	0.746223	0.714975	0.696441	0.782769	0.714975	0	1	0	2.59%	19.86%
11	1.06783	0.876174	0.961522	0.924949	0.915102	0.936295	0.924949	0	1	0	1.06%	15.45%
12	1.08478	0.889971	0.941967	0.906345	0.906345	0.916925	0.906345	0	1	0	0.00%	19.69%
13	1.06275	0.945736	0.993226	0.985713	0.961366	1.00434	0.985713	0	1	0	2.47%	7.82%
14	1.02942	0.862939	0.87832	0.891494	0.87832	0.980809	0.87832	1	0	0	0.00%	17.20%
15	1.13806	0.834793	0.875046	0.88432	0.865314	0.912639	0.875046	1	0	0	1.11%	30.06%
16	1.12213	0.876638	0.946338	0.933805	0.91731	1.00139	0.933805	0	1	0	1.77%	20.17%
17	0.994982	0.871918	0.947107	0.898479	0.898479	0.898479	0.898479	0	1	1	0.00%	10.74%
18	1.03097	0.771097	0.873768	0.838278	0.816722	0.838278	0.838278	0	1	1	2.57%	22.99%
19	0.958223	0.785906	0.846498	0.839849	0.833133	0.865949	0.839849	0	1	0	0.80%	14.09%
20	1.0878	0.868806	0.97094	0.900374	0.900374	0.938122	0.900374	0	1	0	0.00%	20.82%
							Totals	2	18	4	0.88%	17.50%

TABLE 1b. 20 Trials for $n = 6$

Trial	OPASTAT	LOWBD	OPADYN	GREEDY	OPTIMAL	HICOST	Min of 3	OPADYN?	GREEDY	HICOST	DEVN %	SAVEOPA
1	1.38192	1.05927	1.20419	1.15723	1.13812	1.18465	1.15723	0	1	0	1.65%	19.42%
2	1.26817	1.03782	1.16885	1.10995	1.10995	1.17079	1.10995	0	1	0	0.00%	14.25%
3	1.30766	0.937226	1.14352	1.03424	1.03424	1.08734	1.03424	0	1	0	0.00%	26.44%
4	1.39028	1.02838	1.20023	1.09611	1.07332	1.09611	1.09611	0	1	1	2.08%	26.84%
5	1.42086	1.18485	1.21274	1.19612	1.19612	1.23151	1.19612	0	1	0	0.00%	18.79%
6	1.33969	0.927279	1.09828	1.05066	1.00632	1.06633	1.05066	0	1	0	4.22%	27.51%
7	1.20036	0.791894	0.996377	0.874511	0.846262	0.918299	0.874511	0	1	0	3.23%	37.26%
8	1.36207	1.08858	1.25318	1.17622	1.14493	1.2048	1.17622	0	1	0	2.66%	15.80%
9	1.337	1.11319	1.21445	1.1634	1.13994	1.20937	1.1634	0	1	0	2.02%	14.92%
10	1.28349	1.06524	1.227	1.14204	1.1127	1.16834	1.14204	0	1	0	2.57%	12.39%
11	1.20985	0.972564	1.13945	1.03949	1.03949	1.12086	1.03949	0	1	0	0.00%	16.39%
12	1.24912	1.02921	1.13099	1.11009	1.11009	1.16593	1.11009	0	1	0	0.00%	12.52%
13	1.30373	0.920046	1.06812	1.01119	1.01119	1.09683	1.01119	0	1	0	0.00%	28.93%
14	1.32391	1.02181	1.17718	1.06894	1.06894	1.10605	1.06894	0	1	0	0.00%	23.85%
15	1.2761	1.1612	1.25539	1.20498	1.20012	1.28379	1.20498	0	1	0	0.40%	5.90%
16	1.34799	1.07088	1.23371	1.09982	1.09982	1.18642	1.09982	0	1	0	0.00%	22.56%
17	1.25994	1.08613	1.14346	1.17433	1.12628	1.21456	1.14346	1	0	0	1.50%	10.19%
18	1.38227	1.06419	1.21706	1.13063	1.11623	1.14226	1.13063	0	1	0	1.27%	22.26%
19	1.36657	1.05549	1.22818	1.13114	1.12201	1.19751	1.13114	0	1	0	0.81%	20.81%
20	1.40259	1.19367	1.19367	1.11522	1.09783	1.15471	1.11522	0	1	0	1.56%	25.77%
							Totals	1	19	1	1.20%	20.14%

TABLE 1c. 10 Trials for $n = 7$

Trial	*OPASTAT*	*LOWBD*	*OPADYN*	*GREEDY*	*OPTIMAL*	*HICOST*	*Min of 3*	*OPADYN?*	*GREEDY*	*HICOST*	*DEVN %*	*SAVEOPA*
1	1.49311	1.2072	1.39306	1.31786	1.28743	1.31786	1.31786	0	1	1	2.31%	13.30%
2	1.51693	1.13459	1.34871	1.28011	1.23988	1.35619	1.28011	0	1	0	3.14%	18.50%
3	1.60145	1.19883	1.38036	1.33479	1.29185	1.39431	1.33479	0	1	0	3.22%	19.98%
4	1.59725	1.28127	1.43565	1.36337	1.34633	1.40006	1.36337	0	1	0	1.25%	17.15%
5	1.41293	0.995835	1.16834	1.1399	1.06204	1.14506	1.1399	0	1	0	6.83%	23.95%
6	1.37959	1.18874	1.30621	1.23	1.23	1.41485	1.23	0	1	0	0.00%	12.16%
7	1.56199	1.25948	1.40654	1.33181	1.33074	1.37596	1.33181	0	1	0	0.08%	17.28%
8	1.54869	1.24002	1.39145	1.35874	1.32283	1.38964	1.35874	0	1	0	2.64%	13.98%
9	1.54335	1.20525	1.43265	1.26911	1.26911	1.40082	1.26911	0	1	0	0.00%	21.61%
10	1.44116	1.08035	1.28633	1.21775	1.17684	1.24105	1.21775	0	1	0	3.36%	18.35%
							Totals	0	10	1	2.28%	17.63%

TABLE 1d. 10 Trials for $n = 8$

Trial	*OPASTAT*	*LOWBD*	*OPADYN*	*GREEDY*	*OPTIMAL*	*HICOST*	*Min of 3*	*OPADYN?*	*GREEDY*	*HICOST*	*DEVN %*	*SAVEOPA*
1	1.67447	1.37738	1.59107	1.57621	1.49173	1.62136	1.57621	0	1	0	5.36%	6.23%
2	1.8893	1.30812	1.57881	1.45892	1.43041	1.53471	1.45892	0	1	0	1.95%	29.50%
3	1.75605	1.45849	1.64501	1.60143	1.56067	1.65423	1.60143	0	1	0	2.55%	9.66%
4	1.62357	1.21609	1.46007	1.41728	1.36089	1.45426	1.41728	0	1	0	3.98%	14.56%
5	1.88776	1.48836	1.68623	1.60414	1.57632	1.69101	1.60414	0	1	0	1.73%	17.68%
6	1.79422	1.3568	1.58539	1.51112	1.48097	1.54388	1.51112	0	1	0	2.00%	18.73%
7	1.76181	1.36466	1.63913	1.57264	1.51093	1.61807	1.57264	0	1	0	3.92%	12.03%
8	1.78788	1.33919	1.61793	1.53177	1.46294	1.57198	1.53177	0	1	0	4.49%	16.72%
9	1.70453	1.45283	1.57945	1.5726	1.55332	1.61701	1.5726	0	1	0	1.23%	8.39%
10	1.77534	1.40913	1.59422	1.55351	1.52563	1.59797	1.55351	0	1	0	1.79%	14.28%
							Totals	0	10	0	2.90%	14.78%
10 Trials for $n = 9$												
1	1.97885	1.50699	1.7828	1.72631	1.67432	1.77504	1.72631	0	1	0	3.01%	14.63%
2	2.01727	1.55136	1.796	1.7655	1.68764	1.90826	1.7655	0	1	0	4.41%	14.26%
3	1.9714	1.45345	1.70415	1.65426	1.58914	1.74061	1.65426	0	1	0	3.94%	19.17%
4	2.09747	1.61053	1.89353	1.78176	1.73022	1.81716	1.78176	0	1	1	2.89%	17.72%
5	2.04965	1.56562	1.81362	1.75176	1.68928	1.75176	1.75176	0	1	1	3.57%	17.01%
6	2.00108	1.44453	1.76257	1.6862	1.61841	1.69482	1.6862	0	1	0	4.02%	18.67%
7	2.0874	1.66954	1.91653	1.82755	1.80656	1.85814	1.82755	0	1	0	1.15%	14.22%
8	2.03249	1.54577	1.79617	1.71786	1.69215	1.79507	1.71786	0	1	0	1.50%	18.32%
9	1.96806	1.53024	1.74773	1.75354	1.66959	1.79837	1.74773	1	0	0	4.47%	12.61%
10	2.03184	1.61533	1.86611	1.78912	1.74121	1.83791	1.78912	0	1	0	2.68%	13.57%
							Totals	1	9	1	3.16%	16.02%
10 Trials for $n = 10$												
1	2.30239	1.6635	2.0336	1.94962	1.872543	2.04528	1.94962	0	1	0	3.96%	18.09%
2	2.24565	1.79187	2.01723	2.01539	1.96014	2.08531	2.01539	1	1	0	2.74%	11.43%
3	2.3637	1.57991	2.01192	1.89382	1.79713	1.90865	1.89382	0	1	0	5.11%	24.81%
4	2.13992	1.7481	2.03629	1.98261	1.92399	2.04768	1.98261	0	1	0	2.96%	7.93%
5	2.24608	1.64555	2.02902	1.88729	1.87447	2.074	1.88729	0	1	0	0.68%	19.01%

TABLE 1d. Continued

Trial	*OPASTAT*	*LOWBD*	*OPADYN*	*GREEDY*	*OPTIMAL*	*HICOST*	*Min of 3*	*OPADYN?*	*GREEDY*	*HICOST*	*DEVN %*	*SAVEOPA*
6	2.32026	1.82401	2.13392	2.01376	1.96906	2.07533	2.01376	0	1	0	2.22%	15.22%
7	2.19595	1.79397	2.06198	2.0138	1.95038	2.08705	2.0138	0	1	0	3.15%	9.05%
8	2.13871	1.73	2.03842	1.95246	1.89653	2.05759	1.9546	0	1	0	2.86%	9.54%
9	2.27821	1.6967	2.01925	1.96951	1.89607	1.98891	1.96951	0	1	0	3.73%	15.67%
10	2.223	1.82538	2.0915	2.0478	1.99639	2.10834	2.0478	0	1	0	2.51%	8.56%
							Totals	0	10	0	2.99%	13.93%

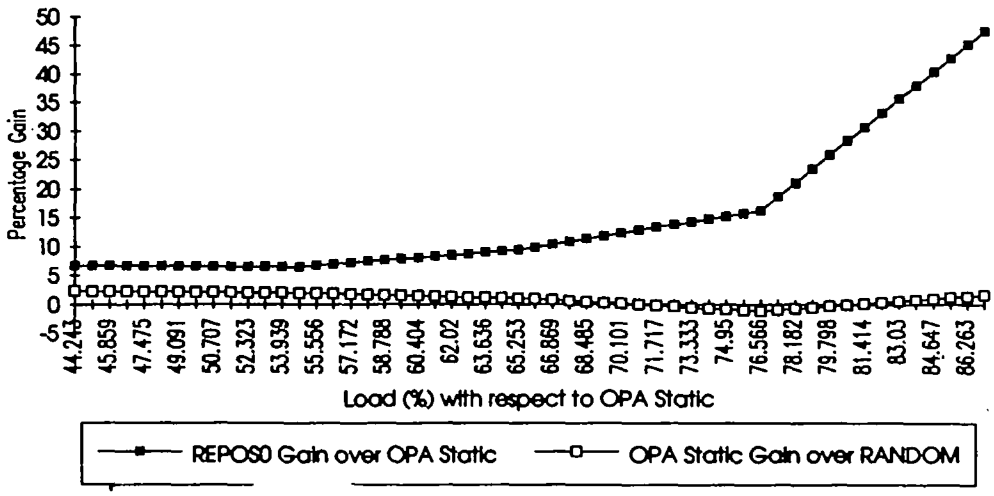

FIGURE 6. Performance of OPA static versus REPOS0 & RANDOM, number of locations = 10.

location $i+1$ and then to location $i+2$. Even such a provision is unavailable today. We assume therefore that once a command to reposition has been given, it has to be executed and only then subsequent commands can be taken up for execution. We also assume that requests are served in the first come first served (FCFS) order. We assume that the arrival process of requests can be modelled by using a Poisson process. Given that a large number of users will be making infrequent requests to a tertiary storage system, this is a reasonable assumption to make (for example see Gnedenko and Kovalenko, 1989). We assume that the cassette request probabilities are such that 70% of the requests are for 10% of the cassettes, 20% for 20% of cassettes and the remaining 10% for 70% of the cassettes. If files that are rarely accessed are put into tertiary storage, then based on practical experience this assumption would be valid.

We assume that the original model for cassette requests holds, i.e. requests are independent of one another and that the request probabilities do not change over time. We investigate the case when there is only one read head. Given these assumptions we ask: does repositioning yield any benefits? How much benefit do we get from using an OPA arrangement over a random arrangement of cassettes?

We conducted several simulation experiments to answer these questions and their results are summarized in Tables 2a–d. In these experiments we varied the load on the system from 30% to 95% and the size of the carousel from 10 to 40 cassettes in steps of 10. The request probabilities were assigned as follows: given $n = 10, 20, 30, 40$ break up n into high load class i.e. 10% of n, medium load class, i.e. 20% of n and low load class consisting of 70% of the n cassettes. We then assign request probabilities uniformly within each class. As an example, if $n = 10$, the request probabilities were set to be: 0.7, 0.1, 0.1, 0.0142857, 0.0142857, 0.0142857, 0.0142857, 0.0142857, 0.0142857 and 0.0142857. The reading time from the tape was assumed to be uniformly distributed over $[0, 2]$ and the rotational time between two adjacent locations to be 0.1. In Table 2a we show the simulation results for the case $n = 10$. The cassettes are arranged as per OPA in all except the columns labeled RANDOM. Note that the average service time will depend on the arrangement of cassettes as well as the repositioning policy. In STATIC OPA, we do not allow

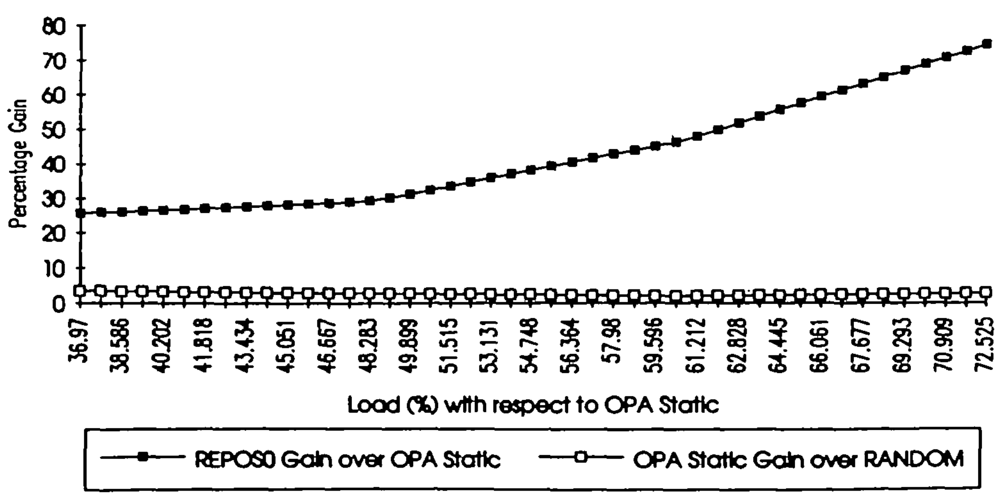

FIGURE 7. Performance of OPA static versus REPOS0 & RANDOM, number of locations = 20.

TABLE 2a. Number of tape locations = 10

	(Static OPA) Static OPA rule			(REPOS) Reposition always		(REPOS[1]) Repositon when 1		(REPOSØ) Reposition empty		(RANDOM: 20 Trials) Tapes randomly placed				
Arrival Rate	Mean Service Time	Load on System (%)	Mean Number in System (1)	Mean Service Time	Mean Number in System	Mean Service Time	Mean Number in System	Mean Service Time	Mean Number in System (2)	Mean Number in System (3)	Max mean in 20 Trials	Min mean in 20 Trials	Gain RESPOSO over Static OPA (1)/(2)−1	Gain Static OPA over Random (3)/(1)−1
0.4	1.094	43.76	0.608	1.107	0.621	1.105	0.615	1.073	0.594	0.649	0.693	0.601	0.023569	0.067434
0.5	1.094	54.7	0.882	1.107	0.907	1.104	0.895	1.077	0.865	0.939	0.997	0.89	0.019653	0.064626
0.6	1.094	65.64	1.29	1.106	1.354	1.102	1.315	1.081	1.277	1.413	1.537	1.308	0.01018	0.095349
0.7	1.094	76.58	2.077	1.106	2.186	1.099	2.132	1.085	2.102	2.416	2.798	2.054	−0.01189	0.163216
0.8	1.094	87.52	4.054	1.105	4.309	1.097	4.026	1.09	3.984	6.055	7.861	4.326	0.01757	0.493587
0.9	1.095	98.55	37.867	1.105	88.948	1.095	36.871	1.094	37.264	118471	308196	2481.3	0.016182	3127.608

TABLE 2b. Number of tape locations = 20

	(Static OPA) Static OPA rule			(REPOS) Reposition always		(REPOS[1]) Repositon when 1		(REPOSØ) Reposition empty		(RANDOM: 20 Trials) Tapes randomly placed				
Arrival Rate	Mean Service Time	Load on System (%)	Mean Number in System (1)	Mean Service Time	Mean Number in System	Mean Service Time	Mean Number in System	Mean Service Time	Mean Number in System (2)	Mean Number in System (3)	Max mean in 20 Trials	Min mean in 20 Trials	Gain RESPOSO over Static OPA (1)/(2)−1	Gain Static OPA over Random (3)/(1)−1
0.3	1.211	36.33	0.466	1.281	0.51	1.274	0.501	1.17	0.45	0.585	0.637	0.514	0.035556	0.255365
0.4	1.213	48.52	0.722	1.282	0.804	1.27	0.774	1.182	0.702	0.933	1.061	0.785	0.02849	0.292244
0.5	1.214	60.7	1.103	1.282	1.269	1.261	1.166	1.189	1.081	1.614	1.994	1.291	0.020352	0.463282
0.6	1.212	72.72	1.79	1.281	2.2	1.249	1.879	1.195	1.741	3.127	5.634	2.208	0.028145	0.746927
0.7	1.21	84.7	3.296	1.279	4.967	1.233	3.385	1.2	3.272	6087.1	17167	7.09	0.007335	1845.814
0.8	1.212	96.96	18.297	1.277	7924.8	1.217	17.71	1.21	17.738	11827.1	1570000	49820	0.031514	6473.887

TABLE 2c. Number of tape locations = 30

	(Static OPA) Static OPA rule			(REPOS) Reposition always		(REPOS[1]) Repositon when 1		(REPOSØ) Reposition empty		(RANDOM: 20 Trials) Tapes randomly placed				
Arrival Rate	Mean Service Time	Load on System (%)	Mean Number in System (1)	Mean Service Time	Mean Number in System	Mean Service Time	Mean Number in System	Mean Service Time	Mean Number in System (2)	Mean Number in System (3)	Max mean in 20 Trials	Min mean in 20 Trials	Gain RESPOSO over Static OPA (1)/(2)−1	Gain Static OPA over Random (3)/(1)−1
0.3	1.325	39.75	0.551	1.419	0.615	1.406	0.593	1.264	0.518	0.751	0.869	0.639	0.063707	0.362976
0.4	1.325	53	0.862	1.412	0.981	1.393	0.927	1.28	0.819	1.345	1.529	1.053	0.052503	0.560325
0.5	1.32	66	1.37	1.407	1.674	1.374	1.467	1.288	1.332	3.095	3.926	2.084	0.028529	1.259124
0.6	1.321	79.26	2.468	1.408	3.448	1.358	2.567	1.302	2.427	4333.8	11991.8	4.981	0.016893	1754.997
0.7	1.323	92.61	7.494	1.405	36.2	1.337	7.738	1.317	7.389	760222	1498000	21673	0.01421	101443.1
0.73	1.323	96.579	16	1.406	8421.1	1.33	16.04	1.32	15.6	2570000	3630000	1580000	0.025641	160624

TABLE 2d. Tape locations = 20

		(Static OPA) Static OPA rule		(REPOS) Reposition always		(REPOS^1) Repositon when 1		(REPOSØ) Reposition empty		(RANDOM: 20 Trials) Tapes randomly placed				
Arrival Rate	Mean Service Time	Load on System (%)	Mean Number in System (1)	Mean Service Time	Mean Number in System	Mean Service Time	Mean Number in System	Mean Service Time	Mean Number in System (2)	Mean Number in System (3)	Max mean % in 20 Trials	Min mean % in 20 Trials	Gain RESPOSO over Static OPA (1)/(2)−1 (%)	Gain Static OPA over Random (3)/(1)−1 (%)
0.2	1.452	29.04	0.357	1.582	0.406	1.577	0.396	1.355	0.327	0.5	0.538	0.41	0.091743	0.40056
0.3	1.43	42.9	0.609	·1.554	0.707	1.535	0.67	1.359	0.569	1	1.105	0.833	0.070299	0.642036
0.4	1.427	57.08	1.003	1.556	1.215	1.518	1.095	1.375	0.97	2.01	2.593	1.347	0.034021	1.003988
0.5	1.434	71.7	1.81	1.563	2.517	1.502	1.956	1.402	1.957	20	98.79	3.566	−0.07511	10.04972
0.6	1.431	85.86	3.929	1.559	9.41	1.468	4.159	1.415	3.947	320230	650807	18406	−0.00456	81503.2
0.66	1.433	94.578	10.568	1.557	8332.9	1.448	11.113	1.428	10.922	1827000	2947000	737994	−0.03241	172879.4

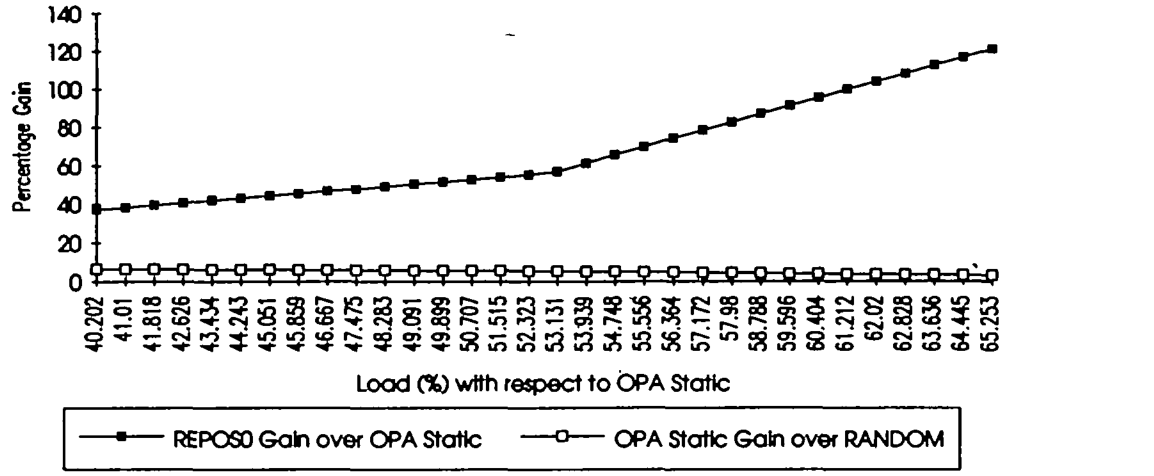

FIGURE 8. Performance of OPA static versus REPOS0 & RANDOM, number of locations = 30.

repositioning and the system load is computed with reference to this policy. In REPOS we always reposition, which as seen from the table adds tremendous overhead to the mean service time because the carousel always rotates back to the cassette with the highest request probability. In REPOS1, we reposition when there is one request and give the repositioning command when that cassette is taken up for service. That is at the time of taking up for service the (lone) request we add a command for repositioning, which is always executed. This adds decreasing amount of overhead as the load increases, as seen from Table 2a. In REPOS0 we reposition when the system is empty. This adds overheads once in a while, but that overhead is not shown in the table as it was difficult to adjust the simulator to compute the overhead added to the request that arrived while the carousel was being repositioned. However note that the mean service time increases with the load because less and less repositioning gets done as the load increases. (If repositioning is not done frequently we are almost back to the STATIC OPA case, and on top of that suffer a penalty whenever repositioning interferes with the next request.) This demonstrates the trade off between repositioning overhead and the reduction in the mean service time. Finally, we generated 20 different random arrangements of cassettes and assumed that no repositioning was done in these 20 cases. We show the results from these arrangements under RANDOM, and give the mean, maximum and minimum of the average number of requests over the 20 cases. For each policy/arrangement we give the average number of requests in the system over a suitable length of simulation. The simulation runs were for 10000 to 9000000 time units. The average number of requests in the system was collected for each time interval of 1000 units to get a standard error for the average number in the system over the entire run. The run length of the simulations was adjusted to keep this standard error within 1–2% of the average number in the system for the entire length of simulation. The standard errors are not shown in the Tables. The simulation was coded in f77 and run on a network of SUN workstations at the Leonard N. Stern school of Business, New York University.

The results show that REPOS is not a good strategy. REPOS0 is the best strategy we have discovered. But even REPOS0 gave at the most 9% improvement over OPA STATIC and that too at low loads. This is shown in the graphs in Figures 6–9. The remarkable fact from

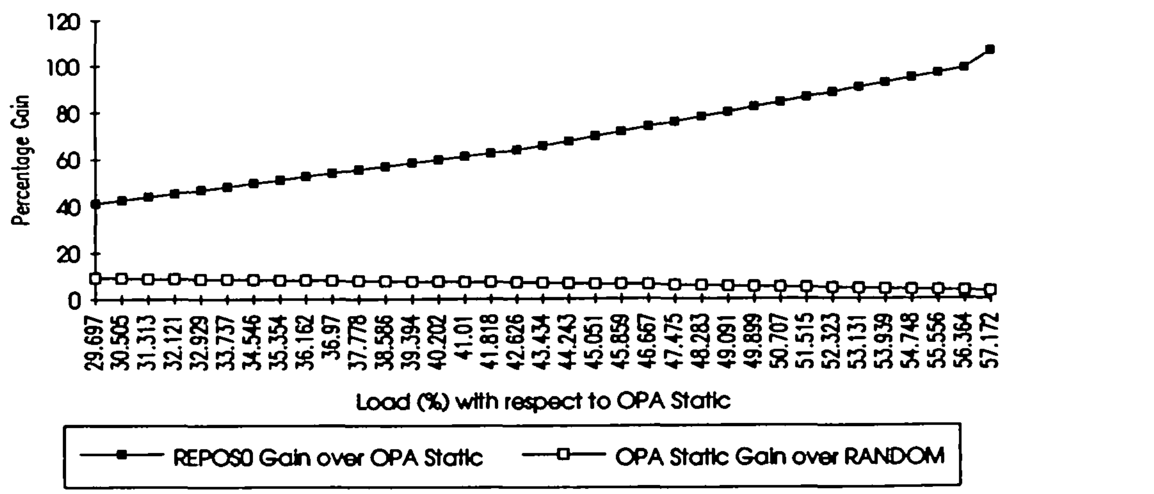

FIGURE 9. Performance of OPA static versus REPOS0 & RANDOM, number of locations = 40.

these simulations is that OPA STATIC outperforms RANDOM significantly and as the carousel size increases, the gain in performance becomes better and better.

The explanation for the improvement in performance over RANDOM, lies almost entirely in the fact that the use of OPA reduces the average service time. The extent of the reduction in service time will depend on the size of the carousel, because RANDOM has a greater chance to deviate from the optimal arrangement as the size increases. Another factor that influences the percentage reduction in the average service time is the ratio between the average time to read from a cassette and the average travel time to move to the cassette location. This ratio would get progressively small as the size of the carousel increases. Thus the relative improvement in the average service time will increase (by using OPA over RANDOM).

5. CONCLUSIONS

In this paper we studied organization schemes of cartridges on a carousel type robotic tape library. The following Table summarizes the results:

	Anticipatory	*Non-Anticipatory*
One head/Single File	Proposition 2	Fujimoto (1991), Proposition 1
Two heads/Single File	Proposition 5	Proposition 6 (even number of locations)
One head/Multiple Files	Proposition 4	Proposition 3
Two heads/Multiple Files	Proposition 5	Proposition 6 (even number of locations)

In addition, we showed that by using the concept of stochastic ordering, all the above results can also be extended to the queueing environment with FCFS policy.

Some questions raised by this work are:

(*i*) *Varying file sizes*. In the file allocation problem we assumed that the files are of equal size. If the file sizes are not the same, we will have packing limitation based on file size and cartridge capacity. This leads to an NP-complete problem as shown in Wong83. Based on the results in this paper a good heuristic can be constructed by first ranking files as per the request probability per unit of size and ordering them in OPA based on these modified probabilities.

(*ii*) *Other queueing disciplines*. We have ignored scheduling problems in proposition 7 by assuming a FCFS service discipline. When file sizes vary, there is some advantage in attending to requests that need smaller files on a priority basis.

(*iii*) *Analysis of other robotic devices*. The carousel is only one type of robotic device, other architectures include cabinets with multiple shelves such as the EXB-120 (by Exabyte Corporation) where the robotic arm picks cartridges from the shelves and places them in up to four parallel drives. Optimal arrangements of cartridges in such architectures and efficient mount schedules of the parallel drives are of interest.

ACKNOWLEDGEMENTS

The authors would like to thank the anonymous referee for his valuable comments which led to the simulations reported in Section 4.

This work was supported by the Applied Mathematical Sciences Research program of the Office of Energy Research, US Department of Energy under Contract DE-AC03-76SF00098.

REFERENCES

Bergmans, P. P. (1973) Minimizing expected travel time on geometrical patterns by probability rearrangements. *Information and Control*, **20**, 331–350.

Fujimoto, A. (1991) *Models for determining storage retrieval strategies for carousel storage and retrieval systems*. Masters Thesis in Management Science, University of Waterloo, Waterloo, Ontario, Canada.

Gnedenko, B. V. and Kovalenko, I. N. (1989) *Introduction to Queueing Theory*, 2nd edn. Birkhauser, Boston, MA.

Grossman, D. D. and Silverman, H. F. (1973) Placement of records on a secondary storage device to minimize access time. *JACM*, **20**, 429–438.

Hardy, G., Littlewood, J. E. and Polya, G. P. (1991) *Inequalities*. Cambridge University Press, Cambridge.

Hofri, M. (1983) Should the two-headed disk be greedy?—Yes, it should. *Inf. Process. Lett.*, **16**, 645–653.

Karp, R. M., McCellar, A. C. and Wong, C. K. (1975) Near-optimal solutions to a 2-dimensional placement problem. *SIAM J. Computing*, **4**, 271–286.

King, R. P. (1990) Disk arm movement in anticipation of future requests. *ACM Trans. Computer Systems*, **8**, 214–229.

Lim, W. K., Bartholdi, J. J. and Platzman, L. K. (1985) Storage schemes for carousel conveyors under real time control. *Material Handling Research Center Technical Report, MHRC-TR-85-10*, Georgia Institute of Technology, Atlanta, GA.

Ranade, S. (1992) *Jukebox and Robotic Libraries*. Meckler, Westport, CT.

Stoyan, D. (1983) *Comparison Methods for Queues and Other Stochastic Models*. John Wiley, New York, NY.

Wolff, R. W. (1989) *Stochastic Modeling and the Theory of Queues*. Prentice-Hall, Englewood Cliffs, NJ.

Yue, P. C. and Wong, C. K. (1973) On the optimality of the probability ranking scheme in storage applications. *JACM*, **20**, 624–633.