Performance of Disk Systems with Two Read/write Heads per Surface

Athena Vakali

Department of Informatics
Aristotle University
54006 Thessaloniki, Greece
email: vakali@eng.auth.gr

Abstract

Our research topic is the performance of two-headed disk systems. Several scheduling algorithms have been adopted to serve read and write requests, and the expected seek has been calculated and compared to that of single-headed disk systems. Data placement schemes have been also studied in conjunction with the scheduling algorithms in order to study the efficiency and fault tolerance of two-headed disk systems. Probability theory and simulation models have been used to achieve performance results and reach conclusions.

Performance of Disk Systems with Two Read/write Heads per Surface

Abstract

Our research topic is the performance of two-headed disk systems. Several scheduling algorithms have been adopted to serve read and write requests, and the expected seek has been calculated and compared to that of single-headed disk systems. Data placement schemes have been also studied in conjunction with the scheduling algorithms in order to study the efficiency and fault tolerance of two-headed disk systems. Probability theory and simulation models have been used to achieve performance results and reach conclusions.

1 Literature Survey

The most common device for secondary storage is the moving arm disk system. In particular, read/write (r/w) heads are mounted on the arm, which moves so that the r/w heads approach several parallel co-centric recording circular surfaces in order to satisfy requests. Surfaces are divided in tracks, whereas the set of tracks with the same radius forms a cylinder. Conventional disks have one head per surface, which serves all the tracks of a surface. All the heads move together to the requested cylinder, but only one head is able to read or write at a time.

Disks were originally designed for file permanent storage, so the primary design criteria were: cost, size and speed. A r/w operation may transfer one or many sectors of stored data. Given the arrangement of surfaces and heads, the time required for a particular operation (e.g. to access, read and transmit a sector of data) involves five components:
(A) select the appropriate surface,
(B) move the appropriate head to the appropriate cylinder (*seek time*),
(C) wait for the required sector to rotate around the location of the r/w head (*latency time*),
(D) read the bytes from the disk surface (*block transfer time*),
(E) transmit bytes to the CPU.
Operation (A) is purely electronic and for this reason the time involved is negligible. Operation (D) depends primarily on the amount of data read, whereas operation (E) depends on the busyness of the channel to the CPU. Thus, operations (B) and (C) have been studied in our research carried out so far.

As an attempt to improve time performance, during the last decade disks with two heads per surface have been commercially available (e.g. Burroughs FD120, DEC RA81, IBM 3380, Sperry Univac 8450). Thus, there is a growing interest in studying the effectiveness and performance gain over cost in these systems. In the aforementioned commercial systems, the two heads/surface remain always in a fixed distance. However, systems with two heads/surface able to move independent of each other in all of the data band have been also proposed and examined. The present research deals with two-headed disks with either fixed or independent heads, introducing new scheduling algorithms and comparing different two-headed models together with the conventional one-headed systems.

Every Input/Output (I/O) device has a queue of pending requests. If the desired disk and controller are available, the request can be serviced immediately, whereas any additional request will need to be queued. Thus, when a request is satisfied, a new request from the queue is picked to be serviced. Scheduling of cylinder requests is a major issue that has been studied extensively in the past. Several scheduling algorithms have been introduced in order to improve seek performance. First-Come-First-Served (FCFS), Shortest-Seek-Time-First (SSTF) and SCAN scheduling algorithms have been examined in many archival papers with certain conclusions. Results have been presented for each one of the algorithms individually as well as for different algorithms compared with each other. Concerning one-headed disk systems, SCAN algorithms and its variations are frequently used for request servicing as in [16, 44], where efficient algorithms evaluate mean waiting time. FCFS and SSTF have been studied and compared in [25, 45], whereas a continuum on disk scheduling algorithms is presented in [22]. Properties of disk scheduling policies in multiprogrammed systems are summarized in [43], where both analytic and simulation results are obtained in order to discuss performance criteria. Various queuing models have been introduced in [17, 23, 31] in most of which Markov chains are used to depict the system states.

Concerning two-headed disk systems, scheduling algorithms have been also studied in various papers. Two derivatives of the SSTF scheduling policy are introduced in [40], the FCFS policy has been used in [3, 4] for request servicing, whereas in [37] the SCAN policy has been applied in a two-headed disk system. The fixed distance separating the two heads is another major point of research interest. References [3, 6, 37, 40] search for the optimal head separation distance, i.e. for the optimum number of cylinders separating the two heads/surface. Under the FCFS policy the optimum separation distance is 0.44657 of the total number of cylinders [3], but under the SCAN policy this distance has been proved to be $C/2 - 1$ ($\lceil C/2 - 1 \rceil$ and $\lfloor C/2 - 1 \rfloor$), when the total number of cylinders C is even (odd).

Another important issue affecting the involved seek as well the latency time is the way of placing data among cylinders. Data placement techniques in conventional disks has been examined in [24, 28, 46], which discuss minimization of r/w head movement when either probability of accessing data records or synchronized disk interleaving is used as storage architecture. The "organ-pipe" arrangement is proven to be optimal when consecutive seeks are independent. The idea of multidisk file design is proposed in [11, 10, 12, 20] for answering partial match queries with minimum cost. Several data allocation methods (Disk Modulo, Symbolic Gray Code, Error Correcting Codes) have been examined in [9, 13, 14, 15, 21, 34] to facilitate partial match retrieval on multiple disks.

Concerning two-headed disk systems, two mathematical models with optimal directory placement have been studied in [5], a technique of reordering cylinders achieves expected seek reduction in [7], while an adaptive data placement scheme balances computer work load in [27]. In [33] a declustering method is used which spreads each file across several disks while [38] proves the "camel-arrangement" to be optimal in two-headed disks. Concerning two-headed disks greedy algorithms are proposed in [26, 29], where a disk arm moves to specific cylinder positions in anticipation of future requests. Recently there has been a considerable interest in parallel replicated disks, where all disks are identical images and store exactly the same data. In [1, 35] analytic models have been developed to evaluate seek performance and certain

expressions have been derived for reads and writes as functions of the number of disk copies of the system.

As described above, research regarding either one- or two-headed disks covers the main topics of Scheduling of requests, Data placement among cylinder disks and Probability measures in specific model development. The research presented here, includes methods and algorithms for each of these topics as described in the following section.

2 Contributions

Our research that has been carried out so far resulted in the following list of publications and reports:
(a) Y.Manolopoulos and A.Vakali: "Seek distances in disks with two independent heads per surface", *Information Processing Letters*, Vol.35, pp.37-42, Jan.1991.
(b) Y.Manolopoulos and A.Vakali: "Request scheduling in disks with two r/w heads per surface", *3rd PanHellenic Conference of Informatics*, Vol.1, pp.171-182, May 1991.
(c) Y.Manolopoulos and A.Vakali: "Expected Seeks in Mirrored Disks", *2nd Hellenic-European Conference on Mathematics and Informatics, HERMIS 94*, Sep.1994.
(d) Y.Manolopoulos and A.Vakali: "Parallel data paths in two-headed disk systems", *submitted for journal publication, Dec.1994.*
(e) Y.Manolopoulos and A.Vakali: "An exact analysis on expected seeks in shadowed disks", *submitted for journal publication, Jan.1995* .
(f) Y.Manolopoulos and A.Vakali: "On partial match retrieval in two-headed disk systems, *submitted for conference presentation, Mar.1995.*

In reference (a) estimates for the average seek distance in a disk system with two r/w heads have been derived, for two different models of heads: randomly positioned or intelligently moved in anticipation of future requests. In this paper a linear storage medium consisting of C cylinders on top of which two r/w heads move independently, is considered. The assumptions made are that cylinders are hit equiprobably under a uniform distribution and both heads may lie on top of any cylinder with equal probability. The mean seek distance traveled is equal to the ratio of the sum of seek distances divided by the number of all the possible combinations of head positions and hit cylinders. According to the first model, the closest head moves to the requested cylinder in order to satisfy the request. Calculations are carried out and the average seek distance for answering a request is approximated by $5C/24$. According to the second model, the heads move by obeying the rule that if the first head points to a cylinder $A < C/2$, then the other head is positioned on top of cylinder $A + \lfloor 2\ (C - A)/3 \rfloor$, otherwise if $A \geq C/2$, then the second head is positioned on top of the cylinder $\lfloor A/3 \rfloor$. Similar calculations carried out for this model approximate average seek distance by $5C/36$. A considerable gain has been shown here in comparison with one-headed disk system, where the mean head movement was approximated by $C/3$.

In reference (b) an optimal policy of request scheduling is applied as a special case of the "traveling salesman" problem. A two-headed disk system of C cylinders is considered and different request scheduling algorithms are adopted for the two heads moving always in fixed

distance. The assumption here is that the two heads remain in fixed distance $d = C/2$, whereas the initial positions of the two heads are on top of the 1st and $(d+1)$-th cylinders. Each head can serve only one of the $[1, d]$ or $[d+1, C]$ intervals. Each request refers to a set of cylinders bounded by a left cylinder (L) and a right cylinder (R). Servicing of a request starts as soon as the previous request is completely served. We categorize the cost due to seeks in two derivatives: the answering cost and the transition cost. The answering cost is measured by the number of cylinders passed from L to R, while the transition cost is measured by the number of cylinders that have to be scanned in order to reach a request's cylinder L after cylinder R of the previous request has been reached. CSCAN algorithms is used for request servicing and three mutually excluding cases have been studied. The first case deals with one head being able to serve all requests, i.e. all L and R of the requests belong to the same interval. An $O(n\ logn)$ algorithm is used to derive optimal scheduling, which reduces the movement cost whereas the answering cost remains the same. The second case deals with the instance that both heads are needed to satisfy the request. Under this case the movement cost is increased, since CSCAN is used with heads remaining in fixed distance. The last case has to do with requests which belong to either of the two previous cases. In the latter case an optimal algorithm is suggested in order to minimize the movement cost. For each of the three cases, analytic formulae have been derived for the answering cost and the arm movement cost.

The performance of two different disk models is examined in reference (c). Models studied here refer to parallel systems with replicated disks which store identical data. Reading data is satisfied by accessing any of the disks since they all store exactly the same data while writing new information must be satisfied by all disk drives in order for them to remain identical copies of each other. Seek time is approximated by the average number of cylinders traveled by the heads when the arm moves from the current cylinder to the requested one and a uniform distribution of requested cylinders is assumed. The first model considers a disk with one r/w head in a replicated mirrored system of only two disk drives. Seek is calculated for reads and writes by the use of a Markov chain with state space the number of distinct head positions over the data band. Calculations are based on a new perspective for the transition probability for the state 1. The results derived indicate a performance improvement in writes varying from 4% to 25% as the read ratio varies from 0.9 to 0.1, while reads remain as optimistic as in [35]. The second model deals with a set of k two-headed disk drives which are all exact copies of each other. The two r/w heads are supposed to move autonomously in all of the data band. Two models are studied, one for independent seeks and the other for dependent ones. In case of independent seeks, the seek distance gain due to the usage of two-headed disks is an almost 45% (respectively, 38%) decrease for reads (respectively, writes). Seeks being dependent results from the fact of all heads being in an identical position when servicing a write request. A Markov chain is used for all possible fluctuations of the number of different cylinder positions. A considerable gain is pointed out which for reads is in the range of 40-50%, whereas for writes is of 40-55%, as the read ratio varies from 0.9 to 0.1. So, the fact of having independent head movement is fully exploited and results seem to boost I/O performance.

Report (d) deals with the expected seek performance of parallel disks having two r/w heads per disk surface. A parallel system of k identical disk drives is considered. Two models are studied, one with autonomous r/w heads and the other with the two heads moving always in a fixed distance. Performance of both models is examined under two different perspectives,

distance. The assumption here is that the two heads remain in fixed distance $d = C/2$, whereas the initial positions of the two heads are on top of the 1st and $(d+1)$-th cylinders. Each head can serve only one of the $[1, d]$ or $[d+1, C]$ intervals. Each request refers to a set of cylinders bounded by a left cylinder (L) and a right cylinder (R). Servicing of a request starts as soon as the previous request is completely served. We categorize the cost due to seeks in two derivatives: the answering cost and the transition cost. The answering cost is measured by the number of cylinders passed from L to R, while the transition cost is measured by the number of cylinders that have to be scanned in order to reach a request's cylinder L after cylinder R of the previous request has been reached. CSCAN algorithms is used for request servicing and three mutually excluding cases have been studied. The first case deals with one head being able to serve all requests, i.e. all L and R of the requests belong to the same interval. An $O(n\ logn)$ algorithm is used to derive optimal scheduling, which reduces the movement cost whereas the answering cost remains the same. The second case deals with the instance that both heads are needed to satisfy the request. Under this case the movement cost is increased, since CSCAN is used with heads remaining in fixed distance. The last case has to do with requests which belong to either of the two previous cases. In the latter case an optimal algorithm is suggested in order to minimize the movement cost. For each of the three cases, analytic formulae have been derived for the answering cost and the arm movement cost.

The performance of two different disk models is examined in reference (c). Models studied here refer to parallel systems with replicated disks which store identical data. Reading data is satisfied by accessing any of the disks since they all store exactly the same data while writing new information must be satisfied by all disk drives in order for them to remain identical copies of each other. Seek time is approximated by the average number of cylinders traveled by the heads when the arm moves from the current cylinder to the requested one and a uniform distribution of requested cylinders is assumed. The first model considers a disk with one r/w head in a replicated mirrored system of only two disk drives. Seek is calculated for reads and writes by the use of a Markov chain with state space the number of distinct head positions over the data band. Calculations are based on a new perspective for the transition probability for the state 1. The results derived indicate a performance improvement in writes varying from 1% to 25% as the read ratio varies from 0.9 to 0.1, while reads remain as optimistic as in [35]. The second model deals with a set of k two-headed disk drives which are all exact copies of each other. The two r/w heads are supposed to move autonomously in all of the data band. Two models are studied, one for independent seeks and the other for dependent ones. In case of independent seeks, the seek distance gain due to the usage of two-headed disks is an almost 45% (respectively, 38%) decrease for reads (respectively, writes). Seeks being dependent results from the fact of all heads being in an identical position when servicing a write request. A Markov chain is used for all possible fluctuations of the number of different cylinder positions. A considerable gain is pointed out which for reads is in the range of 40-50%, whereas for writes is of 40-55%, as the read ratio varies from 0.9 to 0.1. So, the fact of having independent head movement is fully exploited and results seem to boost I/O performance.

Report (d) deals with the expected seek performance of parallel disks having two r/w heads per disk surface. A parallel system of k identical disk drives is considered. Two models are studied, one with autonomous r/w heads and the other with the two heads moving always in a fixed distance. Performance of both models is examined under two different perspectives,

successive seeks being either independent or dependent. Seek time is approximated as in (c) and the choice of the disk drive to serve reading is made by applying the "minimum distance" policy. FCFS policy is used for request servicing and again uniform distribution of requested cylinder positions is assumed. In case of independent seeks, the autonomous heads model shows a 32% to 40% improvement in the total expected seek distance, whereas the fixed headed model shows an almost 50% improvement over the one-headed model and a 15% to 20% improvement over the autonomous headed model. In case of successive seeks being dependent a Markov chain is used to depict the actual number of different cylinder positions. The autonomous headed model shows a 43% to 60% performance improvement compared to one-headed disk systems, i.e. this model seems to behave more than twice as well as the one-headed disk model. In case of fixed headed model there is a 50% improvement in total expected seek over the one-headed model, while there is a 15% gain in writes and almost a 15% worsening on reads performance compared to the autonomous headed model. As shown in Table 1, taking advantage of dependency between successive seeks causes total expected seek to improve at almost 50% (read ratio $r = 0.1$) over the non-dependant scheme when servicing writes while the improvement rate is not as high when servicing reads.

	Independent seeks			Dependent Seeks		
	r=0.9	r=0.5	r=0.1	r=0.9	r=0.5	r=0.1
k=2	0.13	0.2	0.29	0.1	0.16	0.2
5	0.08	0.23	0.41	0.08	0.17	0.2
10	0.07	0.26	0.45	0.08	0.17	0.2

Table 1: Total expected seek as percentage of $C = 200$, in autonomous two-headed disk.

An earlier approximate analytic model ([35]) for a shadowed disk system with k disks is revisited in report (e) under a different perspective which takes into account the disk scheduling policy applied (nearer server rule) and the total number of cylinders (C) per disk. The model here is a one-headed disk system with shadowed disks on which reading data is satisfied by accessing any of the disks, while writing new information must be satisfied by all disks since they all are identical copies of each other. Exact results are derived since earlier results of [35] were approximate. In general, a uniform distribution of requested cylinders is assumed. Again, a Markov chain is used to describe the new process, with each state i meaning that there are exactly i distinct cylinder positions occupied by the k r/w heads. The calculation of the transition function is based on probability measures based on cylinder subintervals on which r/w head moves. Subintervals are determined by the number of distinct cylinder positions occupied by the k r/w heads. The presented model is proven to behave slightly better in case of writes with an improvement almost less than 1% and the two models converge as the value of total number of cylinders C increases. In case of reads the model of [35] is too optimistic with respect to the exact values for the expected seek distance traveled and the two models converge as the value of C increase.

Partial match retrieval is a recent topic on which several research work has started the results of which are presented in (f). A two-headed disk system with C cylinders (C an even number here without loss of generality) is considered having the two heads in fixed distance of

$C/2$ with request involving N distinct cylinders out of the C ones. Requests are served by a SCAN scheduling policy. Some declustering method is applied and intraparallelism is examined in a two-headed system in order to satisfy a partial match query. Reduction of distance traveled as a consequence of having heads at two different cylinders simultaneously is calculated here in terms of the number of arm stops. A probability distribution function is derived for the number of stops that the moving arm mechanism makes on top of hit cylinders, i.e. the expected value of arm stops is calculated. When compared to a single headed disk system, it is shown that the gain in stops is linear on the number of cylinder requests and theoretically may reach 50% on the average.

3 Future Research

Our current research deals with the partial match retrieval problem as described in the previous section, in association with specific declustering and data allocation methods. A simulative study is the next step in order to verify the results that appeared in the analytic study, has been described in report (f) of the previous section. The research having been carried out already should be completed by the examination of several different declustering policies, which would yield specific results for expected seek times when answering partial match queries. We believe, that based on the simulation results certain conclusions can be drawn for two-headed systems by comparing the different declustering algorithms between them as well as with the conventional one-headed disk systems. Information placement among parallel disks is a major research theme, so the behavior of specific allocation schemes like Error Correcting Codes, Disk Modulo and Gray Codes will be sufficient in order to clarify whether a certain scheme will be ideal for two-headed systems or not. These methods as well as introducing new algorithms suitable for two-headed systems will be studied for reading and writing servicing of requests in general.

In addition, the case presented in reference (b) can also be analyzed with requests referring to a set of successive cylinders bounded by the (Left, Right) cylinders. All of these models should be studied for both cases of two-headed systems, i.e. heads moving independently and heads moving always in a fixed distance.

Future research should also expand in new scheduling policies. FCFS and SCAN algorithms have been used already in references (a), (b), and (d), whereas similar research can be carried out for SSTF algorithm as well as for variations of SCAN (e.g. CSCAN, VSCAN) algorithm. Simulation must be carried out for this case too and comparisons between different two-headed models as well as with single-headed models can be presented. Research conducted in reports (c), (d) and (e) has used Markov chains to perform calculations and new probability measures could be found also for the new algorithms developed. Expected seek was studied in all of (a), (b), (c), (d), (e), and (f) research reports. Latency time must be studied too since it affects considerably the overall performance. So, future research results should present specific measures for both expected seek and latency time and study the advantages and the possible drawbacks of introducing the new models.

Results in all of the research topics described in this section together with the results that

have been obtained so far as described in the previous section, will complete the aim of the research plans.

4 Conclusions

Research conducted so far has examined specific scheduling algorithms (SCAN, FCFS) in certain analytic models of two-headed disk systems which have either autonomous r/w heads or heads remaining always in a fixed distance. The requests refer to cylinder positions and Markov chains have been used in some of the models in order to calculate expected seek distance. Conclusions made so far prove that two-headed models improve performance in general and they are more reliable specially for the case of parallel disk drives which have also been studied. Partial match retrieval is the topic on which current research elaborates in order to minimize expected seek by the use of special data clustering schemes of files among disk cylinders. The research should expand in the area of declustering and data allocation, where new models need to be introduced together with the scheduling policies. Other query types (except partial match queries) should also be examined and comparisons have to be made in order to emphasize the efficiency of the new models. Simulation should support the theoretical conclusions of the new analytic models.

We believe that the conclusions of this research may be useful and hopefully applicable, in order to improve disk systems performance and make decisions as of whether the cost of the second head is overpassed by the overall systems performance. The documented results and conclusions for two-headed models will be critical for adopting disks with two r/w heads per disk surface.

References

[1] D. Bitton and J. Gray: Disk Shadowing, *Proceedings of the 14th VLDB Conference*, pp.331-338, LA, CA, 1988.

[2] F.W. Burton and J.G. Kollias: Optimizing Disk Head Movement in Secondary Key Retrievals, *The Computer Journal*, Vol.22, No.3, pp.206-208, 1979.

[3] A. Calderbank, E. Coffman and L. Flatto: Optimum Head Separation in a Disk with Two r/w Heads, *Journal of the ACM*, Vol.31, No.4, pp.433-466, 1984.

[4] A. Calderbank, E. Coffman and L. Flatto: Sequencing Problems in Two-server Systems, *Mathematics of Operational Research*, Vol.10, pp.585-598, Nov 1985.

[5] A. Calderbank, E. Coffman and L. Flatto: Optimal Directory Placement on Disk Storage Devices, *Journal of the ACM*, Vol.35, No.2, pp.433-446, 1988.

[6] A. Calderbank, E. Coffman and L. Flatto: A Note Extending the Analysis of Two-head Disk Systems to more General Seek Time Characteristics, *IEEE Transactions on Computers*, Vol.38, No.11, pp.1584-1586, 1989.

[7] S.D. Carson and P. Reynolds: Adaptive Disk Reorganization, *Computer Science Report Series of University of Maryland*, Jan 1989.

[8] S.D. Carson and P. Vongsathorn: Error bounds on Disk Arrangement using Frequency Information, *Information Processing Letters*, Vol.31, pp.209-213, 1989.

[9] M.Y. Chan: A Note on Redundant Disk Modulo Allocation, *Information Processing Letters*, Vol.20, pp.121-123, 1985.

[10] M.Y. Chan and W.L. Chung: Optimal Multi-disk Partial Match File Designs, *Information Processing Letters*, Vol.28, pp.149-155, 1988.

[11] C.C. Chang and C.Y. Chen: Performance of Two-disk Partition Data Allocation, *BIT*, Vol.27, pp.306-314, 1987.

[12] C.C. Chang and C.Y. Chen: A Note on Allocating k-ary Multiple Key Hashing Files among Multiple Disks, *Information Sciences*, Vol.55, pp.69-76, 1991.

[13] C.C. Chang, H.Y. Chen and C.Y. Chen: Symbolic Gray Code as a Data Allocation Scheme for Two-disk Systems, *The Computer Journal*, Vol.35, No.3, pp.299-305, 1992.

[14] C.C. Chang, H.F. Lin and C.Y. Chen: Optimality Properties of the Disk Modulo Allocation Method for k-ary Multiple Key Hashing Files, *The Computer Journal*, Vol.35, No.4, pp.390-394, 1992.

[15] C.Y. Chen and C.C. Chang: On GDM Allocation Method for Partial Range Queries, *Information Systems*, Vol.17, No.5, pp.381-394, 1992.

[16] E.G.Jr. Coffman and M. Hofri: On the Expected Performance of Scanning Disks, *SIAM Journal of Computing*, Vol.11, No.1, pp.60-70, 1982.

[17] E.G.Jr. Coffman and M. Hofri: Queuing Models on Secondary Storage Devices, *Purdue University*, update of the identically named chapter in "Stochastic Analysis of Computer and Communications Systems", H.Takagi, Elsevier Science, Publ. BV 1990.

[18] H.C. Du: Disk Allocation Methods for Binary Cartesian Product Files, *BIT*, Vol.26, pp.138-147, 1986.

[19] H.C. Du and J.S. Sobolewski: Disk Allocation for Cartesian Product Files on Multiple Disk Systems, *ACM Transactions on Database Systems*, Vol.7, pp.82-101, 1982.

[20] C. Faloutsos : Multiattribute Hashing using Gray Codes, *ACM*, Vol., No., pp.228-238, 1986.

[21] C. Faloutsos and D. Metaxas: Disk Allocation Methods using Error Correcting Codes, *IEEE Transactions on Computers*, Vol.40, No.8, pp.907-914, 1991.

[22] R. Geist and S. Daniel: A Continuum on Disk Scheduling Algorithms, *ACM Transactions on Computing Systems*, Vol.5, No.1, pp.77-92, 1987.

[23] C.C. Gotlieb and G.H. MacEwen: Performance of Movable-Head Disk Storage Devices, *Journal of the ACM*, Vol.20, No.4, pp.604-623, 1973.

[24] D.D. Grossman and F.H. Silverman: Placement of records on a Secondary Storage Device to Minimize Access Time, *Journal of the ACM*, Vol.20, No.3, pp.429-438, 1973.

[25] M. Hofri: Disk scheduling: FCFS vs. SSTF Revisited, *Communications of the ACM*, Vol.23, No.11, pp.645-653, 1980.

[26] M. Hofri: Should the Two-headed Discs be Greedy? Yes, it Should, *Information Processing Letters*, Vol.16, pp.83-86, 1983.

[27] K.A. Hua and C.Lee: An Adaptive Data Placement Scheme for Parallel Database Computer Systems,*Proceedings of the 16th VLDB Conference*, pp.493-506, 1990.

[28] M.Y. Kim: Synchronized Disk Interleaving, *IEEE Transactions on Computers*, Vol.35, No.11, Nov 1986.

[29] R.P. King: Disk Arm Movement in Anticipation of Future Requests, *ACM Transactions on Computer Systems*, Vol.8, No.3, pp.214-229, 1990.

[30] J.G. Kollias: An Estimate of Seek Time for Batched Searching of Random or Indexed Sequential Files, *The Computer Journal*, Vol.21, No.2, pp.132-133, 1978.

[31] C.H.C. Leung: Analysis of Disk Fragmentation using Markov Chains, *The Computer Journal*, Vol.26, No.2, pp.113-116, 1983.

[32] C.H.C. Leung: A Model for Disk Locality Referencing, *The Computer Journal*, Vol.26, No.3, pp.196-198, 1983.

[33] M. Livny, S. Khoshafian and H. Boral: Multi-disk Management Algorithms, *Proceedings of the ACM SIGMETRICS Conference*, pp.69-77, 1987.

[34] J. Li, J. Srivastana and D. Rotern: A Multi-dimensional Declustering Method for Parallel Database Systems, *Proceedings of the 18th VLDB conference*, Vancouver, CA. 1992.

[35] R.W. Lo and N.S. Matloff: Probabilistic Limit on the Virtual Size of Replicated Disc Systems, *IEEE Transactions on Knowledge and Data Engineering*, Vol.4, No.1, pp.99-102, 1992.

[36] Y. Manolopoulos and J.G.Kollias: Estimating Disk Head Movement in Batched Searching, *BIT*, Vol.28, pp.27-36, 1988.

[37] Y. Manolopoulos and J.G.Kollias: Performance of a Two-headed Disk System when Servicing Database Queries under the SCAN policy, *ACM Transactions of Database Systems*, Vol.14, No.3, pp.425-442, 1989.

[38] Y. Manolopoulos and J.G.Kollias: Optimal Data Placement in Two-headed Disk Systems, *BIT*, Vol.30, pp.216-219, 1990.

[39] Y. Manolopoulos: Probability Distributions for Seek Time Evaluation, *Information Sciences*, Vol.60, No.1-2, pp.29-40, 1991.

[40] I.P. Page and R.T. Wood: Empirical Analysis of a Moving Disk Head Model with two Heads Separated by a Fixed Number of Tracks, *The Computer Journal*, Vol.24, No.4, pp.339-341, 1981.

[41] R.L. Rivest: Partial Match Retrieval Algorithms, *SIAM Journal of Computing*, Vol.5, No.1, pp.19-50, 1976.

[42] D. Rotem: Analysis of Disk Arm Movement for Large Sequential Reads, *Proceedings of the 11th ACM PODS Conference*, pp.47-54, San Diego, CA 1992.

[43] T.J. Teory and T.B. Pinkerton: A Comparative Analysis of Disk Scheduling Policies, *Communications of the ACM*, Vol.15, No.3, pp.177-184, 1972.

[44] C.C. Wang and B.P. Weems: Some considerations in using VSCAN Disk Scheduling with Optimal Page Arrangements, *Proceedings of the ACM SIGMETRICS 88 Conference*, pp.273, 1988.

[45] N.C. Wilehelm: An Anomaly in Disk Scheduling: a Comparison of FCFS and SSTF Seek Scheduling using an Empirical Model for Disk Access, *Communications of the ACM*, Vol.19, No.1, pp.13-17, 1976.

[46] C.K. Wong: Minimizing Expected Head Movement in One-dimensional and Two-dimensional Mass Storage System, *ACM Computing Surveys*, Vol.12, No.2, pp.167-178,1980.