Dynamic file allocation in disk arrays

Report**Author(s):**

Zabback, Peter; Weikum, Gerhard; Scheuermann, Peter I.

Publication date:

1990

Permanent link:

<https://doi.org/10.3929/ethz-a-000571235>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

ETH, Eidgenössische Technische Hochschule Zürich, Departement Informatik, Institut für Informationssysteme 147

Eidgenössische
Technische Hochschule
Zürich

Departement Informatik
Institut für
Informationssysteme

Gerhard Weikum
Peter Zabback
Peter Scheuermann

Dynamic File Allocation in Disk Arrays

December 1990

Eidg. Techn. Hochschule Zürich
Informatikbibliothek
ETH-Zentrum
CH-8092 Zürich

Author's address:

Informationssysteme
ETH-Zentrum
CH-8092 Zurich, Switzerland

e.mail: weikum@inf.ethz.ch

© 1990 Departement Informatik, ETH Zürich

Dynamic File Allocation in Disk Arrays

Gerhard Weikum, Peter Zabback, and Peter Scheuermann*

ETH Zurich

Department of Computer Science

Information Systems – Databases

CH-8092 Zurich, Switzerland

Abstract

Large arrays of small disks are being considered as a promising approach to high performance I/O architectures. In this paper we deal with the problem of data placement in such a disk array. The prevalent approach is to decluster large files across a number of disks so as to minimize the access time to a file and balance the I/O load across the disks. The data placement problem entails determining the number of disks and the set of disks across which a file is declustered. Unlike previous work, this paper does not assume that all files are allocated at the same time but rather considers dynamic file creations and expansions. This makes the placement problem considerably harder because each placement decision has to take into account the current allocation state and the access frequencies of the disks and the existing files. As a result, both file creation and expansion may involve partial reorganization on one or more disks. The paper proposes heuristic algorithms for the placement of dynamically created files and file expansions. The algorithms provide a good compromise between maximizing I/O performance of the disk array and minimizing the work invested in partial reorganizations. The paper presents preliminary performance results of various alternative algorithms under a synthetic workload.

1 Introduction

To meet the requirements of high performance I/O architectures, disk arrays [35] are being considered as a promising approach. In a disk array, a large number of small disks are used rather than relatively few large disks, and a very high bandwidth interconnect (i.e., hundreds of Megabytes per second) is used for the data transfer between disks and memory. Such an architecture not only provides a higher I/O bandwidth, but has also lower costs compared to large disks. The high bandwidth is achieved because a disk array has many arms so that more I/Os can be processed in parallel. This I/O parallelism can be exploited in two different ways:

1. A higher number of independent I/Os can be performed in parallel. While this does not significantly improve the response time of small I/Os, it can improve the I/O throughput drastically.

*Permanent address: Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, Illinois 60208-3118. This work was performed while the author was visiting ETH Zurich.

2. The response time of large I/Os (i.e., I/Os that request many consecutive blocks) can be decreased by orders of magnitude. Suppose, for example, an entire relation is read into memory in order to compute a join. If the requested relation is distributed over multiple disks, then all fragments of the relation can be read in parallel. This sort of distribution is also known as *declustering* or *striping*.

In this paper, we deal with software-controlled disk arrays, also known as independent drive disk arrays [23], disk farms, or simply multi-disk systems [26]. In this type of disk array, multiple disks can cooperate on a single data request while, at the same time, other requests can be directed to a single disk. The file system is responsible for the distribution of the data and the disk load balancing. Thus, the placement of files can be tailored to the application. For example, the response time of large I/Os such as reading an entire file can be minimized by carefully choosing the number of disks and the set of disks across which a file is declustered. We believe that file systems and database systems, which often have highly skewed data access distributions, crucially need the flexibility of software-controlled disk arrays for load balancing.

1.1 The Data Placement Problem

Data placement within a disk array is a crucial issue for both applications with mostly small I/Os and applications with large I/Os. For small I/Os, load balancing is the critical issue. For large I/Os, minimizing the access time and load balancing are equally important objectives. In order to exploit the advantages of a disk array, we have to address the following issues.

- *Type of declustering unit:*

Declustering can be based on relations (i.e., sets of records) or files. In the first case, a relation is partitioned into fragments, e.g., based on key ranges, and the fragments are allocated on different disks. In the second case, a file is partitioned into runs of consecutive blocks, and these runs are allocated on different disks.

In this paper, we consider only the second case for the following reason. In applications that deal with complex objects such as office documents [51], it may be desirable to decluster a single object that is stored in a large set of blocks. This case can be dealt with by handling the large complex object as a separate file (not necessarily in the sense of a separate OS file but rather as a separate set of blocks that constitute a "storage cluster" [1],[4],[5],[16],[18],[20],[22],[25],[39]). On the other hand, for relations that consist of many small records, the mapping of records into the blocks of the underlying files may be important to reflect a sort order of records. This mapping can be defined in an additional design step on top of file declustering. If a relation is stored physically sorted in the blocks of the underlying file(s), then key-range queries, for example, translate into byte-range accesses which can be parallelized by file declustering. Using files as the declustering unit promises to deal with different sorts of objects in a uniform way, and it is potentially useful for both database systems and OS file systems.

- *Degree of declustering:*

Determining the number of disks across which a relation or file is declustered is a critical performance issue. Small files (e.g., smaller than a track) are obviously better off if they are entirely allocated on one disk, since this allows reading or writing the entire file or any portion of it in a single set-oriented I/O [49]. For large files, a high degree of declustering increases the I/O parallelism, but it also increases the maximum seek time and maximum rotational delay among the disks that serve a read or write request. In addition, the CPU costs of processing a request may increase with the degree of declustering as more device driver work has to be performed. Of course, these disadvantages do not apply to small I/Os that read or write only a single block of the file. So the optimal degree of declustering of a file depends on both the size of the file and its access characteristics.

- *Load balancing:*

When we have decided to decluster a relation or file across w disks, and if w is smaller than the number of disks in the disk array, we still have the problem of selecting the w disks on which the relation or file is actually placed. An intelligent selection is crucially important for disk load balancing. This can be achieved by taking into account statistics (or estimates) of the access frequencies of disks and files. In practice, disk load balancing is achieved by hand tuning the data placement periodically [26]. Automating data placement would be an important step toward self-tuning systems, saving an enormous amount of system administrator work.

- *Dynamic allocation and reorganization:*

The data placement problem is already hard if one assumes that all relations or files are allocated at the same time. In practice, however, this is an unrealistic assumption. File systems are highly dynamic, and many database system applications require the dynamic creation (or deletion) of relations. Moreover, during the evolution of an application, relations or files may grow quite a bit. The resulting dynamic data placement problem may in turn require *partial reorganizations* on heavily utilized disks. Finally, long-term changes of access frequencies may require *global reorganizations* on the entire disk array.

This paper addresses the *dynamic data placement* problem for files. This entails determining the optimal degree of declustering and selecting the set of disks across which a file is declustered. Our main contributions are the following:

- Given an average request size (number of blocks) for a file, and the throughput requirements (number of requests per second) for the disk array, we present a method for deriving the file's "optimal" degree of declustering such that the response time of requests is minimized. This derivation is based on a simple analytical model for response time and throughput as functions of the degree of parallelism in serving a request (see [13] for similar considerations).

- Given a file's degree of declustering, we propose algorithms for selecting the disks across which the file is declustered such that the overall load is balanced across all disks. Because we deal with dynamic file creations and expansions, we consider various sorts of partial reorganizations. The developed algorithms are based on heuristics, since even the static data placement problem is known to be NP-complete [7]. For space limitations, we do not consider global reorganizations in this paper.

Note that our approach to load balancing is independent of how a file's degree of declustering is determined, and what type of declustering unit is used. Thus, our algorithms are also useful for declustering based on relations, in a shared-nothing architecture where the disk load balancing also affects the CPU load balancing, or in disk arrays with workload assumptions that differ from ours.

1.2 Related Work

Declustering (or "striping") data across multiple disks has originally been proposed in [24] and [37]. The performance studies in [26],[31],[36], and [43] show a significant improvement for different sorts of workloads and different sorts of multi-disk architectures. The impact of the degree of declustering on I/O throughput has been studied in [6]. The idea of disk arrays was first introduced in [35]. Disk arrays have motivated the design of novel file systems such as [10],[28],[45], and will presumably influence the file systems of several computer vendors ([17],[29]). However, none of the proposed file systems seems to be flexible enough to cope well with conventional on-line transaction processing as well as accesses to complex objects. The design of our own file system FIVE has been influenced by some ideas introduced in XPRS [45].

Declustering, for performance as well as availability reasons, is employed in several database systems, including commercial systems such as NonStopSQL [47] and Teradata [48] as well as research prototypes such as Gamma [9] and Bubba [3]. These systems use relation fragments as the declustering unit; the partitioning schemes include round-robin, hash partitioning, key-range partitioning, and a hybrid of the latter two ([12],[13]). Other database systems support similar partitioning schemes but require the database administrator to allocate space on different disks manually ([19],[33]).

The data placement problem (i.e., determining across how many disks and across which disks a relation or file is declustered) has been addressed only in the Bubba project [7]. The degree of declustering is chosen based on an analytical model that estimates the performance for a given degree of declustering. The disk selection is based on the notion of *heat*, which is the estimated access frequency of a relation. To achieve good load balancing, disk-resident relations are placed in descending heat order on the disks with the smallest accumulated heat. Note that this heuristic algorithm assumes that all relations are allocated at the same time so that the allocation itself does not pose any problems.

None of the mentioned work considers the dynamic creation of files or relations and the possibly necessary reorganizations. Previous work on reorganization in both file systems (e.g., [38]) and database systems (e.g., [22],[32],[40]) has concentrated on the reorganiza-

tion of a single disk for the purpose of compaction or improving the clustering of records and blocks.

The work presented here aims at automating the data placement and reorganization decisions in a disk-array-based file system. It is most closely related to the data placement work in Bubba [7]. However, it extends this work in that we do not assume that all files are allocated at the same time but rather allow dynamic file creations and expansions.

1.3 Outline of the Paper

The rest of this paper is organized as follows. In Section 2, we give a short description of our prototype file system called FIVE, which serves as the experimental platform for our work. In Section 3, we discuss the performance tradeoffs in dynamic data placement. In Section 4, we develop heuristic data placement algorithms for dynamically created files. In Section 5, we discuss placement heuristics for file expansions. In Section 6, we present preliminary simulation results on dynamic file creations. Finally, we conclude with an outlook on future work on data allocation and reorganization.

2 The File System: FIVE Concepts

We have developed an experimental file system called FIVE that is supposed to exploit the potential advantages of disk arrays. FIVE stands for File System with Adaptive Enhancements. As one might guess, FIVE is based on five concepts, namely blocks, runs, extents, regions, and files.

The smallest unit of data that is managed by FIVE is a *block*. A block is a contiguous fixed-length area on one disk. It is the minimal transfer unit between disk and memory. We require that the block size is the same for all disks and all files. Note that this requirement is not really restrictive because, by using set-oriented I/O, multiple blocks can be transferred in one I/O operation.

A physically contiguous collection of logically consecutive blocks of a file is a *run*. A run consists of the largest number of logically consecutive blocks that reside on one disk. The run size is the "striping granularity" in the sense of ([6],[17]); a number of logically consecutive runs that reside on different disks form a "stripe".

A physically contiguous collection of one or more runs on one disk is called an *extent*. An extent is described by its disk number, its start address, and its size in blocks. A collection of extents located on different disks is called a *region*. All extents of a region have the same run size and the same number of runs per extent. Thus, a region is described by the run size, the extent size, the number of extents, and a list of the start addresses of the extents (i.e., disk number and address of the first block).

Finally, a *file* is stored as a collection of regions. When a file is created, its space is allocated in one region. When a file grows during its lifetime, the file is expanded by allocating additional regions. If the growth factor of a file can be estimated in advance, larger regions can be allocated in order to avoid excessive expansions.

Figure 1: File organization in FIVE

The net effect of this file organization is that runs are placed on w different disks in a round-robin fashion, with those runs that end up on the same disk being clustered together in one extent. The example in Figure 1 illustrates the five introduced concepts. The figure shows the physical allocation of two files and their logical block numbering. Each of the two files consists of one region. The grey region consists of three extents, each of which consists of two runs that are vertically stacked in the figure. The run size of this region is 2 blocks. The black region of the other file consists of two extents, each of which consists of one run of 3 blocks.

The number of extents of a region is called the *width* of the region. The number of blocks of a run is called the *depth* of the region. The number of runs of an extent is called the *height* of the region. Thus, width times depth times height yields the total number of blocks of the region. In the example of Figure 1, the grey region has width 3, depth 2 and height 2; the black region has width 2, depth 3, and height 1. The notion of regions is similar to the organization of "two-dimensional extents" (2D-extents) proposed in XPRS [45]. However, a 2D-extent is always placed on a set of consecutive disks; that is, if the first part is located on disk i , then the rest of the 2D-extent resides on disks $i+1$ through $(i+w-1) \bmod n$ where w is the width of the 2D-extent and n the number of disks in the disk array. We consider this constraint as unnecessarily restrictive, and therefore propose the more flexible notion of a region. In fact, using 2D-extents rather than regions would complicate the data placement problem algorithms that are proposed in this paper.

Logical block numbers are translated into disk block addresses through the *region table* of a file. Figure 2 shows the structure of a region table. We assume that all region tables of frequently used files fit into main memory so that I/Os for region table look-ups are negligible. Note that the separation of logical block numbers and disk block addresses is a prerequisite for all sorts of reorganizations such as relocating extents or combining multiple regions.

In addition to the region tables and a free space bit map for each disk, FIVE keeps the following free space information for each disk d_i :

- *freelist(d_i)*: a list containing the addresses of free areas of disk d_i . The list is kept in decreasing size order (i.e., largest free area first).

Figure 2: Structure of a region table for a file consisting of m regions

- $free(d_i)$: the total number of free blocks on disk d_i . This is not necessarily contiguous space. $free(d_i)$ is the sum of the sizes of all areas in $freelist(d_i)$.
- $maxfree(d_i)$: the largest contiguous area on disk d_i . $maxfree(d_i)$ corresponds to the first entry of $freelist(d_i)$.

The workload of the system is reflected in a metric called *heat* [7], where the heat of an object is the access frequency of the object over some period of time (e.g., one day or week). Only accesses that result in disk I/O are counted so as to factor out the effect of caching. The following heat statistics are (continuously or periodically) collected while the system is running.

- $H(e)$: the heat of an extent e , i.e., the sum of the access frequencies of the extent's blocks.
- $H(d_i)$: the heat of a disk d_i , i.e., the accumulated heat of the extents of the disk.

In addition to heat, the *temperature* of an object is another useful metric, where temperature is defined as follows [7]:

- $T(x) = \frac{H(x)}{size(x)}$: the temperature of an extent, or file x . This metric reflects both the access frequency and the size of an object. Thus, for a large object to have the same temperature as a small object, the large object must be accessed more frequently.

Our file system FIVE maintains the following two sorted lists that are used in the data placement algorithms:

- *heatlist*: a global list of the disks in increasing heat order ($H(d_1) \leq H(d_2) \leq H(d_3) \leq \dots \leq H(d_N)$). This list is also used to calculate the average disk heat avg_heat .
- *templist*(d_i): for each disk d_i , a list containing the addresses of allocated extents on d_i in decreasing temperature order.

3 Performance Tradeoffs

The following performance goals have guided the development of our dynamic data placement algorithms.

- *Minimal access time to consecutive portions of a file,*
i.e., minimizing the response time for read and write requests.
- *Balanced disk load,*
i.e., distributing the load across the disks as uniformly as possible. This is an important goal for both I/O throughput and file access response time because it minimizes queuing delays. Note that disk load balancing is a crucial issue for both shared-nothing systems [44], where the disk load balance directly affects the CPU load balancing, and shared-memory systems, where disk load balancing and CPU load balancing are orthogonal issues.
- *Minimal costs of partial reorganizations,*
i.e., minimizing the extra disk load caused by reorganization steps that result from certain placement decisions. Such placement decisions are in turn a consequence of the other two performance goals. For example, if an extent is placed on a particular disk for load balancing reasons and this disk does not have enough contiguous space, then a partial reorganization is necessary to reclaim sufficient free space.

Figure 3: Incompatible performance goals

Note that we do not aim at tuning the data placement for maximum throughput, i.e., maximizing the throughput at which the disk array saturates, because this would be at the expense of response time. In practice, disks are often underutilized (e.g., with 50 % utilization) on purpose to guarantee acceptable response time (see, e.g., [15],[42]). Throughput requirements are taken into account by limiting the number of disks that serve a single read/write request so that the required request rate can be satisfied (see Section 4.1).

The above three performance goals may be incompatible in some situations, as illustrated in the example shown in Figure 3. The example shows a file f that consists of one region with two extents f_{11} and f_{12} . Suppose that the file is expanded by allocating a new region with two extents f_{21} and f_{22} , and suppose that reading the entire file is the dominating operation. Typically, the file access time would become minimal if the new extents were placed

on disks d_3 and d_4 , which do not yet hold any extents of file f . However, this would even increase the load imbalance between the two cooler disks and the two hotter disks. On the other hand, placing the two new extents on disks d_1 and d_2 sacrifices the goal of optimal file access performance. A similar tradeoff exists between the first two performance goals, on the one hand, and the goal of minimal reorganization costs, on the other hand. The algorithms in Sections 4 and 5 aim at providing a good compromise with respect to our three performance goals.

4 Algorithms for File Creation

In this section, we present algorithms for the dynamic data placement problem that is posed by the creation of new files. The algorithms decide 1) across how many disks a file will be declustered (degree of declustering), and 2) on which disks the file will be placed (disk selection for load balancing). Subsection 4.1 deals with determining the degree of declustering. Subsections 4.2 and 4.3 deal with the disk selection problem, i.e., load balancing. For the sake of modularity, we first present two elementary building blocks for partial reorganizations in Subsection 4.2. These building blocks are used in the description of the disk-selection algorithms that are presented in Subsection 4.3. The performance tradeoffs of alternative algorithms are briefly discussed in Subsection 4.4.

4.1 A Method for Determining the Degree of Declustering

When a new file is created, a single region of size S is allocated. We assume that S is provided by the client of the file system, based on an estimate of the expected file growth. The width w , depth d , and height h of the allocated region are determined as follows, aiming at a minimum access time to the (new) file while satisfying the throughput requirements of the entire disk array.

- Let R be the estimated average request size of the file, i.e., the number of consecutive blocks that are read or written in a single request. Let P be the degree of parallelism in serving such a request, i.e., the number of disks across which the requested blocks are declustered. In the appendix of the paper, we develop a formula for the (single-user) response time of the request as a differentiable function $A(P)$ of P . Solving the equation $\frac{dA}{dP} = 0$ yields the optimum degree of parallelism P_{opt} . With less parallelism, the potential for reducing the transfer time of the request is not fully exploited; with a higher degree of parallelism, the additional gain in transfer time is outweighed by the additional increase of the maximum seek and rotational latency of the involved disks.
- The appendix also contains a formula for the disk array's maximum throughput τ as a function of the overall average request size $\bar{R}$ and the average degree of parallelism $\bar{P}$. Given a required throughput τ^0 , we can solve the inequation $\tau \geq \tau^0$, yielding the maximum average degree of parallelism $\bar{P}_{max}$ as a function of τ^0 and $\bar{R}$.

- Now, our heuristic approach to minimizing response time while observing the throughput requirements is to view $\bar{P}_{max}$ as an upper limit for P_{opt} . Thus, for a file with request size R , we choose the following effective degree of parallelism:

$$P_{eff} = \min (P_{opt}, \bar{P}_{max}, N)$$

where N is the number of disks in the disk array.

- To achieve a degree of parallelism P_{eff} for a request size R , we have to decluster all consecutive portions of size R over P_{eff} different disks, so that P_{eff} runs can be accessed in parallel. This is achieved by choosing an appropriate depth d for the region that is to be allocated, according to the formula

$$d = \left\lceil \frac{R}{P_{eff}} \right\rceil$$

Of course, for larger requests to the same file, the degree of parallelism should be higher. This is achieved by choosing the degree of declustering of the entire file as large as possible. Also, a degree of declustering that is higher than the degree of parallelism of a single request allows parallelism between multiple independent requests to the same file. This consideration yields the following formula for the region width:

$$w = \min \left(N, \left\lceil \frac{S}{d} \right\rceil \right)$$

Finally, we obtain the height of the region according to the formula:

$$h = \left\lceil \frac{S}{w \times d} \right\rceil$$

Note that the above result affects only the degree of declustering of a file. The disk selection algorithms that are described below do not depend on how the values for w , d and h are determined. For example, if we wanted to maximize the saturated disk throughput rather than minimize file access response time, the optimal value of w would probably be smaller, but our disk selection algorithms would still be appropriate for load balancing.

Note that the allocated region of size $w \times d \times h$ may be larger than the requested size S . Such an overallocation is unavoidable if one wants to have all runs and all extents of a region to have the same size. While this property keeps the region table small and makes the addressing of blocks efficient, it has the disadvantage that the internal fragmentation of a file (i.e., $1 - \text{file size} / \text{allocated space}$) can be relatively high. Throughout this paper, however, we consider disk space as an ample resource.

4.2 Building Blocks for Partial Reorganizations

When a dynamic data placement algorithm decides to allocate an extent of a file or region on a particular disk, a partial reorganization may have to be performed on this disk for one of the following two reasons:

- A) There is not enough contiguous space available to satisfy the allocation request.

B) The allocation of the new extent would cause the disk to become too hot (because the extent belongs to a hot file) and may thus adversely affect the load balance in the disk array.

Case A) is addressed by performing a partial disk compaction, and case B) is addressed by performing a reorganization step that we call "disk cooling". These two building blocks are described in the following two subsections.

4.2.1 Partial Disk Compaction (PDC)

A partial disk compaction is necessary when a disk d_i has enough free space for a new extent e but does not have enough contiguous space (i.e., the condition $\maxfree(d_i) < size(e) \leq free(d_i)$ holds). In performing the compaction, the amount of data that is moved should be minimized. Since this is essentially a knapsack problem, we employ a relatively simple heuristic algorithm.

Figure 4 illustrates the effect of the PDC algorithm. In a first step, two addresses low and $high$ determined such that the total size of the free areas between low and $high$ is greater than or equal to $size(e)$ and the total size of the existing extents between low and $high$ is minimized. In the example of Figure 4, we pick $addr(A_1)$ as low and $addr(e_4)$ as $high$ (i.e., extents e_2 and e_3 are moved in order to increase the size of the free area at $addr(A_1)$). After this step, enough contiguous space is available at $addr(A_1)$, and the entries for A_2 and A_3 are removed from $freelist(d_i)$.

Figure 4: Effect of the PDC algorithm

4.2.2 Disk Cooling (DC)

When a disk becomes too hot so that it causes significant load imbalance in the disk array, the heat of the disk can be reduced by moving some of its data to a cooler disk. It is reasonable to move only entire extents so as not to interfere with the goal of minimizing the file access time. For the same reason, an extent is usually not moved to a disk that already holds an extent of the same region (because this would effectively change the width of

the region). If a file consists of multiple regions, one could even exclude disks that already hold an extent of the same file.

Determining the best migration candidates for disk cooling is essentially a knapsack problem, since we want to minimize the amount of data that is to be moved while ensuring that the disk heat will be reduced by a specified amount. A heuristic criterion for selecting the extent(s) to be moved is the temperature $T(e)$ of an extent e . In contrast to the heat metric, the temperature reflects both the benefit and the cost of the reorganization, where the benefit is the achieved decrease of the disk's heat and the cost is proportional to the size of the moved extent(s). A disk cooling algorithm that implements these heuristics is given in Figure 5. The algorithm keeps cooling the disk (i.e., moves extents to cooler disks) until the heat of the disk drops below a given threshold. The algorithm terminates prematurely if the movement of an extent would cause a cooler disk to become hotter than the disk that is to be cooled.

```

Input: disk number  $d_i$ 
      desired heat threshold  $H$ 
repeat
  step1: Determine extent that is to be moved
    Select the extent  $e$  with the highest temperature  $T(e)$ .
  step2: Determine target disk  $d_t$ 
    Select the coolest disk  $d_t$  from the heatlist such that  $d_t$  does not already hold an extent
    of the region to which  $e$  belongs. Exit if moving  $e$  from  $d_i$  to  $d_t$  would cause  $d_t$  to become
    hotter than  $d_i$ .
  step3: Relocate extent
    Move extent  $e$  from  $d_i$  to  $d_t$  (and update the region table of the file to which  $e$  belongs).
    This step may in turn require a partial disk compaction on  $d_i$ .
  step4: Recalculate heat
     $H(d_i) := H(d_i) - H(e)$ 
     $H(d_t) := H(d_t) + H(e)$ 
until  $H(d_i) \leq H$ 

```

Figure 5: Algorithm for disk cooling (DC)

4.3 Disk Selection Algorithms

In this subsection, we present alternative algorithms for selecting the disks across which a newly created file is placed. The alternatives reflect different priorities of the performance goals described in Section 3. In determining the set of disks across which the file is declustered, our algorithms are driven by information about the heat of disks (i.e., the heatlist) and the allocation states of disks as described below. The alternatives use both sorts of information, but with different priorities.

A straightforward approach to selecting w disks is to pick the w coolest disks by looking up the first w entries of the heatlist. This approach, which is adopted from the data placement algorithm of Bubba [7], seems to be a good heuristic for load balancing. However, in a disk array with dynamic file creations, the disk selection also has to take into account

```

Input: region width  $w$ 
       extent size  $e = d \cdot h$  (depth  $d$ , height  $h$ )

result_set := {}; escalation_flag := false
start_label:
for each class  $C$  of allocation states (in preference order) do
  for each eligible  $C$  disk  $d_i$  in heatlist do
    if  $d_i$  not in result_set or escalation flag is set then
      if  $H(d_i) > \text{avg\_heat} + \text{tolerance}$ 
        and there exists a cooler disk  $d_j$  that is not yet in result_set
        then perform algorithm DC ( $d_i, H(d_i) - \text{avg\_heat}$ )
        fi
      case (allocation state of  $d_i$ )
        CS: allocate extent
        ES: execute algorithm PDC ( $d_i, e$ ) and allocate extent
        NES: perform space reclamation algorithm on  $d_i$  and allocate extent
      esac
      result_set := result_set  $\cup \{d_i\}$ 
      if enough extents allocated then exit fi
    fi
  od
od
if not enough extents allocated then
  set escalation flag (i.e., drop the constraint that the allocated extents reside on different disks)
  goto start_label
fi

```

optional
step

Figure 6: Generic disk selection algorithm

the current allocation states of the disks, where the *allocation state of a disk d_i with regard to a new allocation request of size e* is one of the following states:

- CS: contiguous space available, i.e., $e \leq \text{maxfree}(d_i)$
- ES: enough space available, i.e., $\text{maxfree}(d_i) < e \leq \text{free}(d_i)$
- NES: not enough space available, i.e., $e > \text{free}(d_i)$

Obviously, disks with allocation state CS are preferable over ES disks which in turn are preferable over NES disks. Picking an ES disk for an extent allocation requires a partial disk compaction. Picking an NES disk for an extent allocation requires moving some data to a different disk so as to reclaim sufficient free space. Such a *space reclamation* can be accomplished by an algorithm similar to the disk cooling algorithm of Figure 5. Like the DC algorithm, one or more extents have to be moved to one or more cooler disks that do not yet hold an extent of the region or file to which the moved extent belongs. This rule prevents negative effects on load balancing and file access performance. Unlike the DC algorithm, the criterion for selecting the extents that are moved is size (e.g., best-fit) rather than temperature. Also, the termination condition is based on the amount of space that is to be reclaimed rather than the decrease of the disk heat.

Given the additional information about disk allocation states, we obtain various disk selection algorithms by grouping the three possible allocation states (or a subset of them) into classes of equally preferable states and defining a preference order between the classes. Then, a generic disk selection procedure, which is shown in Figure 6, runs as follows. The heatlist is processed in multiple rounds. For each class of states, in preference order, we make a pass over the heatlist and select all eligible disks that are in this class. The allocation states within the same class do not affect whether a disk is selected or not, but they may require different reorganization steps on the selected disk. Note that each round may change the allocation state of some disks. The disk selection procedure terminates when it has selected the required number of disks. If, due to unusual space allocation conditions, not enough disks could be selected after all rounds, then we drop the constraint that the extents of a region must reside on different disks and retry the whole procedure.

1: {CS, ES, NES}	All disks are selected according to the heatlist.
2: {CS, ES}	NES disks are not selected.
3: {CS}	ES or NES disks are not selected.
4: {CS} < {ES}	CS disks are selected in the first round ES disks may be selected in the second round; NES disks are not selected.
5: {CS, ES} < {NES}	CS or ES disks are selected in the first round; NES disks are selected in the second round.
6: {CS} < {ES, NES}	CS disks are selected in the first round; ES or NES disks are selected in the second round.
7: {CS} < {ES} < {NES}	CS disks are selected in the first round; ES disks are selected in the second round; NES disks are selected in the third round.

Figure 7: Possible allocation state classes and preference orders

The above considerations lead to a family of alternative algorithms, one for each possible grouping of allocation states and preference ordering between the resulting classes. The reasonable alternatives are shown in Figure 7. Of the seven possible alternatives, the following three are considered as viable options:

- Alternative 1 – *Heat Balancing (HB)*: The first w entries of the heatlist are selected, regardless of their allocation states.
- Alternative 5 – *Space-restricted Heat Balancing (SHB)*: NES disks are selected only if there are not enough CS or ES disks. This may be reasonable since the space reclamation on an NES disk involves additional disks, whereas a partial disk compaction on an ES disk is a purely local operation.
- Alternative 7 – *Cost Minimization (CM)*: The heatlist is processed in three rounds, aiming at minimal reorganization costs.

In some alternatives, a relatively hot CS disk may be selected whereas a cool ES or NES disk will eventually be not selected. Moreover, allocating the new extent on the hot disk will possibly make it even hotter. In order to avoid large load imbalances that may result from this behavior, we propose the following disk cooling technique. If the heat of the se-

lected disk exceeds the average disk heat in the disk array and if there exists at least one cooler disk that has not yet been selected for the allocation of the new file, then we invoke the disk cooling algorithm of Figure 5 to reduce the heat of the selected disk. That is, one or more extents are removed from the disk and are reallocated on one or more cooler disks (see Section 4.2.2). Of course, this cooling step is reasonable not only for CS disks, but should be applied also to ES disks if they satisfy the above stated conditions. After this step, the new extent is allocated on the selected disk. Note that it is important to perform the disk cooling algorithm before allocating the new extent, because this can save a partial disk compaction if the cooling step turns an ES disk into a CS disk.

4.4 Discussion of Alternatives

In this subsection, we discuss some performance implications of the presented file creation algorithms, using the following example for illustration. Suppose we want to allocate a new file of size $S=16$ blocks on a disk array containing 12 disks with track size 10 blocks. The average request size to this file is $R=10$ blocks. Further suppose that our analytical model for the optimal degree of parallelism for such a request yields $P_{opt} = 5$ and $\bar{P}_{max} = 10$ (assuming disk parameters like the ones used in the appendix, an overall average request size $\bar{R}=10$, and a required throughput of $\tau^0 \approx 50$ requests per second). From this the method described in Section 4.1 derives depth $d=\lceil 10/5 \rceil=2$; width $w=\lceil 16/2 \rceil=8$ as the file's degree of declustering, and height $h=\lceil 16/16 \rceil=1$. That is, a region is allocated consisting of 8 extents each of which consists of one run of 2 contiguous blocks. Now suppose that the heatlist and the current allocation states with regard to a request of size 2 are as shown in Figure 8. Assume that the *avg_heat* of the disks is between $H(d_6)$ and $H(d_7)$. Our three main alternatives for the disk selection would produce the following results:

Figure 8: An example of the *heatlist*

- *HB*: selects the disks $\{d_1, d_2, d_3, d_4, d_5, d_6, d_7, d_8\}$, which causes partial disk compactions on d_1, d_2, d_4, d_5 , and d_7 , and space reclamations on d_3 and d_8 , and, optionally, disk cooling on d_8 .
- *SHB*: selects disks $\{d_1, d_2, d_4, d_5, d_6, d_7, d_9, d_{10}\}$, which causes partial disk compactions on d_1, d_2, d_4, d_5 , and d_7 , and, optionally, disk cooling on d_9, d_{10} .

- *CM*: selects disks $\{d_6, d_9, d_{10}, d_{11}, d_{12}, d_1, d_2, d_4\}$ (in this order), which causes partial disk compactions on d_1, d_2 , and d_4 , and may invoke the disk cooling algorithm on $d_9, d_{10}, d_{11}, d_{12}$.

This example clearly shows the main difference in the performance of the three alternatives. Alternative HB stresses the load balancing goal and is therefore likely to perform more reorganization work. Alternatives SHB and CM, on the other hand, try to avoid partial reorganizations as long as possible and rather sacrifice the load balancing goal.

5 Algorithms for File Expansion

In this section, we discuss algorithms for file expansions, i.e., adding a new region to an existing file. In principle, the allocation of an additional region is not much different from the allocation of a new file. Thus, most of the steps of the file creation algorithm(s) of Section 4 can be directly carried over. However, the following considerations suggest some changes to the disk selection algorithms.

An existing file may already be declustered across many disks; in fact, very large files are likely to be declustered across all disks in the disk array. This fact can influence the disk selection for a new region in two different ways:

1. If there are some disks that do not yet hold any extents of the file, then these disks are preferable over disks on which the file currently resides. To characterize such favorable disks, we introduce an additional *allocation state UU – unused – with regard to the file that is to be expanded*. Selecting UU disks for the newly allocated region increases the I/O parallelism in accesses to the entire file or large portions of it.
2. If there is no UU disk (i.e., the file is already declustered across all disks) or if there are not enough UU disks to allocate the entire new region, then such disks are preferable that have enough contiguous free space in physical proximity to an existing extent that belongs to the same file. Physical proximity here means that the existing extent and the free space are within a small distance of k tracks, where k depends on the disk type characteristics. Disks with this property are characterized by the additional allocation state *CSN – contiguous space nearby – with regard to an allocation request of the file that is to be expanded*. Selecting a CSN disk for a new region is a heuristic for minimizing the seek time when multiple extents are read from or written to the same disk, ideally in a single set-oriented I/O.

The UU state does actually occur in combination with one of the other allocation states, except with CSN. So we end up with the following set of 7 possible allocation states:

{UU-CS, UU-ES, UU-NES, CSN, CS, ES, NES}

Now, we can again define classes of equally preferable allocation states and a preference order between the resulting classes. Then, the generic disk selection procedure of Figure 6 can be easily adapted to handle file expansions. Of course, seven different allocation states give rise to a large number of alternative state class definitions and preference orders. Two particularly intriguing alternatives are the following:

- {UU-CS} < {UU-ES} < {UU-NES} < {CSN} < {CS} < {ES} < {NES}

This alternatives makes up to seven passes over the heatlist. If a file is already declustered across all disks, then the first three state classes are irrelevant. In general, this alternative gives high priority to file access time by preferring UU disks regardless of whether their second state component is CS, ES, or NES.

- {UU-CS} < {CSN} < {CS} < {UU-ES} < {ES} < {UU-NES} < {NES}

This alternative aims at minimal reorganization costs by preferring CSN and CS disks over ES and NES disks regardless of whether they are in state UU or not.

6 Preliminary Performance Measurements

6.1 Description of the Experiment

We have started investigating the performance of the proposed algorithms. Our testbed consists of the implemented FIVE file system, a load generator, and a simulated I/O system that is based on the C-based, process-oriented simulation language CSIM [41]. The simulated disk array consists of 12 (unsynchronized) disks with the parameters shown in Figure 9.

Block Size	4096 Bytes	Average Rotational Latency	8.3 ms
Track Size	10 Blocks	Average Seek Time	16 ms
Tracks per Cylinder	4	Transfer Rate	2.5 MB/s
Cylinders per Disk	624	Cylinder Switch Time	5.0 ms
		Head Switch Time	2.5 ms

Figure 9: Parameters of the simulated disks

In a first experiment, which is described here, we concentrated on dynamic file creations in a synthetic complex-object scenario. The test database consists of approximately 2500 complex objects each of which corresponds to one file. The file size is exponentially distributed with an average of 400 KBytes (i.e., 100 blocks of 4KB each). The simulated workload consists of file creations (5% of all operations), file deletions (2.5%), reads (60%), and writes (32.5%). We assume that a complex object is always accessed entirely; that is, the read/write request size is the file size¹. Files are selected for access according to a 90–10 Zipf-like distribution. That is, 90% of all read/write requests access only 10% of all files. We approximated this skewed access distribution by applying a linear transformation to a normal distribution. The average request size is $\bar{R} = 150$ blocks. Note that there is a correlation between the heat and the size of a file, as larger files cause more data transfer

1. Even though this is a debatable assumption, we have obtained interesting insights into the behavior of our algorithms. In addition, we are not aware of any published analysis of real-life complex-object accesses, which could serve as a counterpart of, for example, the Berkeley trace data analysis of Unix file accesses [34].

work. Thus, since the load generation is driven by the skewed heat distribution of files, the average request size is larger than the average file size.

The degree of declustering of a file is determined according to the method of Section 4.1. We assume that the required throughput is $\tau^0 = 32$ requests per second. From this we derive a maximum degree of declustering of 3, using the formula for $\bar{P}_{max}$ that is given in the appendix.

We have compared the following strategies for dynamic file allocations:

- *Vanilla (V)*: this allocates file extents so as to balance the space utilization of the disks. The V strategy serves as an example of a strategy that disregards load balancing.
- *Cool Vanilla (C-V)*: this may additionally invoke the disk cooling algorithm (DC) of Section 4.2. At each create operation, the heat imbalance of the disks is checked; and if the heat of a disk is higher than 1.1 of the average disk heat in the disk array, then the DC algorithm is invoked.
- *Simple Heat Balancing (HB)*: as described in Section 4.3.
- *Cool Heat Balancing (C-HB)*: like HB with the possibility of DC invocations.
- *Simple Cost Minimization (CM)*: as described in Section 4.3.
- *Cool Cost Minimization (C-CM)*: like CM with the possibility of DC invocations.

We generated test databases by running the operation mix of our synthetic workload, starting with an empty database and terminating at a total space utilization of 75%. Such a database was generated for each of the above six strategies, thus constructing six different databases. In addition, we constructed a seventh database from one of the generated ones by statically reallocating all files in descending heat order, that is, by applying the static data placement algorithm of [7]. The seven test databases are referred to as V-DB, C-V-DB, HB-DB, C-HB-DB, CM-DB, C-CM-DB, and OPT-DB. In the actual measurements, our synthetic workload was run against each of these databases at varying arrival rates of requests. Each run consists of 10000 file-system operations.

6.2 Performance Results

Figure 10 shows the disk heat distribution of the generated test databases. The shown heat metric is the sum of the number of accesses to the blocks of a disk. A low standard deviation across the disks in the disk array is an indicator for good load balance. The real touchstone for load balancing, however, is the standard deviation of the disk utilizations during the actual measurements. These additional figures, at an arrival rate of 30 requests per second, are also shown in Figure 10.

According to the utilization figures of Figure 10, one would expect that the HB and CM strategies clearly outperform the vanilla strategy, with CM being slightly better than HB. We performed measurements by applying each strategy to the test database that was created by the strategy, referred to as x-DB, and by applying all strategies to the V-DB and the OPT-DB. The measured response time of create operations (including an initial write) and

of read/write operations are shown in Figure 11. The main observations from this experiment are summarized in the following.

after the database generation phase							
	V-DB	C-V-DB	HB-DB	C-HB-DB	CM-DB	C-CM-DB	OPT-DB
avg. heat	665221	665221	665221	669097	665659	669558	665221
min. heat	593802	551880	662181	663227	661762	663227	665220
max. heat	781905	780041	675874	713558	675830	719069	665221
standard deviation	55372	64947	4143	14123	3790	15698	1
during the measurement phase							
	V on V-DB	C-V on C-V-DB	HB on HB-DB	C-HB on C-HB-DB	CM on CM-DB	C-CM on C-CM-DB	CM on OPT-DB
min. util.	0.80	0.87	0.81	0.81	0.88	0.88	0.88
max. util.	0.99	0.99	0.97	0.97	0.96	0.96	0.93
avg. util.	0.91	0.93	0.89	0.89	0.92	0.92	0.90
standard deviation	0.06	0.04	0.04	0.04	0.03	0.03	0.01

Figure 10: Heat distribution and load balance for different strategies

- Disk cooling was never invoked in the strategies C-HB and C-CM, since the load balancing of HB and CM is already so good that the imbalance threshold for DC invocations was never exceeded during the measurement phase. Thus, C-HB and C-CM perform exactly like HB and CM, respectively. Therefore, the curves for C-HB and C-CM are omitted in Figure 11. For the vanilla strategy, disk cooling had a beneficial effect; that is, strategy C-V clearly outperformed the simple strategy V. For example, at an arrival rate of 30 requests per second, the response time of strategy V (on the V-DB) was 1.66 times higher than that of strategy C-V (on the C-V-DB). On the OPT-DB, V and C-V performed identically; so the C-V curve is omitted in Figure 11 for this case.
- The vanilla strategies V and C-V both perform well at low arrival rates, as the load imbalance does not yet cause queueing effects. With increasing arrival rate, requests may become queued at a disk. Then, load imbalance causes a dramatic increase of response time because the hottest disk becomes a bottleneck. For example, at an arrival rate of 31 requests per second, the response time of V is almost twice as high as that of the best strategy CM. This observation clearly demonstrates the need for a load balancing strategy.
- At low arrival rates, the HB and CM strategies perform partial disk compactions, without achieving any benefits since request queueing is not yet a problem. This explains why these strategies perform worse than the V strategy which benefits from its space balancing in this case. At higher arrival rates, the reorganizations pay off, for they

Figure 11: Response time of file creations and read/write requests

were actually necessary to improve load balancing and hence decrease the average queue length of the disks.

	V on		HB on			CM on		
	V-DB	OPT-DB	HB-DB	V-DB	OPT-DB	CM-DB	V-DB	OPT-DB
#PDC	-	-	21	49	4	3	8	2
I/Os per PDC	-	-	40	126	40	60	57	55
Blocks per PDC	-	-	1400	4400	1400	2100	1940	1760

Figure 12: Frequency and costs of partial disk compactions during the measurement phase

- Partial disk compaction turned out to be pretty expensive, as it moved up to 4400 blocks. The HB algorithm was far too aggressive, which resulted in excessive PDC invocations and high response time, especially on the badly balanced V-DB. The frequencies and costs of PDC invocations are shown in Figure 12.

On the HB-DB and OPT-DB, the problem is that the HB strategy tends to be overly sensitive in that it prefers an ES disk over a slightly hotter CS disk. This situation arises fairly frequently because the disk load is already relatively well balanced so that a few additional requests may already change the order of disks in the heatlist. Obviously, the HB algorithm needs to be augmented by incorporating a heat difference threshold such that approximately equally hot disks are considered as equally preferable.

- The strategy CM performs best at high arrival rates, as it achieves reasonably good load balance and avoids reorganizations as long as possible.

7 Conclusion

In this paper, we have developed algorithms for the dynamic data placement problem in disk arrays. These algorithms can be used for allocating newly created files and for file expansions with the goals of optimal file access time and disk load balancing at acceptably low reorganization costs. Our algorithms are useful for dealing with both advanced DBMS applications such as office document management where a large complex object corresponds to a file, and conventional relations of small records where file declustering can be exploited by an appropriate mapping of records into blocks (i.e., to reflect a sort order). The simulation results for complex-object accesses are a first step toward gaining quantitative insight into the performance of our algorithms. We are planning on a series of performance experiments under different sorts of workloads.

The partial reorganizations of our algorithms are intended to maintain good file access performance and good disk load balancing over fairly long periods of time. In the long term, the performance may nevertheless deteriorate, for example, because of changing file access frequencies. Therefore, a global reorganization of the entire disk array will probably be necessary once in a while. We are investigating, when and how such a global reorganization should be performed.

An aspect that we have disregarded so far is reliability. With the increasing number of disks in a disk array, the mean time to failure (MTTF) decreases dramatically. Several solutions to this problem have been proposed, either based on error-correcting codes such as parity blocks [11],[14],[30],[35],[46], or based on replicated data [2],[8],[21]. We plan to investigate to what extent our dynamic data placement algorithms need to be extended if the disk array contains replicated data. We expect that hardware-based schemes such as parity striping on a per sector basis fit well with our approach in that they do neither introduce major complications to the disk selection algorithms nor any degradation of load balancing.

We believe that disk arrays will play an important role in future database system architectures. On the other hand, it is likely that data placement in large disk arrays will not be (easily) manageable for most system administrators. Our long-term goal is a high performance database system that automatically adapts itself to the workload and does therefore not require a human system administrator for performance tuning [50]. Developing algorithms that automate data allocation and reorganization in disk arrays is a crucially important (sub-)problem. We believe that this paper is a promising approach toward solving this problem.

References

- [1] Benzaken, V., An Evaluation Model for Clustering Strategies, in the O2 Object-Oriented Database System, *Proceedings of the International Conference on Database Theory*, 1990
- [2] Bitton, D., Gray, J.N., Disk Shadowing, *Proceedings of the 14th International Conference on Very Large Data Bases*, 1988, pp. 331-338
- [3] Boral, H., Alexander, W., Clay, L., Copeland, G., Danforth, S., Franklin, M., Hart, B., Smith, M., Valduriez, P., Prototyping Bubba, A Highly Parallel Database System, *IEEE Transactions on Knowledge and Data Engineering*, Vol. 2, No. 1, 1990, pp. 4-24
- [4] Carey, M.J., DeWitt, D.J., Richardson, J.E., Shekita, E.J., Storage Management for Objects in EXODUS, in: Kim, W., Lochovsky, F.H., (ed.) *Object-Oriented Concepts, Databases, and Applications*, Addison-Wesley, 1989, pp. 341-369
- [5] Chang, E., Katz, R., Exploiting Inheritance and Structure Semantics for Effective Clustering and Buffering in an Object-Oriented DBMS, *Proceedings of the SIGMOD International Conference on Management of Data*, 1989, pp. 348-357
- [6] Chen, P.M., Patterson, D. A., Maximizing Performance in a Striped Disk Array, *Proceedings of the 17th International Symposium on Computer Architecture*, 1990, pp. 322-331
- [7] Copeland, G., Alexander, W., Boughter E., Keller, T., Data Placement in Bubba, *Proceedings of the SIGMOD International Conference on Management of Data*, 1988, pp. 99-108
- [8] Copeland, G., Keller, T., A Comparison of High-Availability Media Recovery Techniques, *Proceedings of the SIGMOD International Conference on Management of Data*, 1989, pp. 98-109
- [9] DeWitt, D.J., Ghandeharizadeh, S., Schneider, D., Bricker, A., Hsiao, H.-I., Rasmusen, R., The Gamma Database Machine Project, *IEEE Transactions on Knowledge and Data Engineering*, Vol. 2, No. 1, 1990, pp. 44-62
- [10] Dibble, P.C., Scott, M.L., Beyond Striping: The Bridge Multiprocessor File System, *ACM Computer Architecture News*, Vol. 17, No. 5, 1989, pp. 32-39
- [11] Garcia-Molina, H., Salem, K., The Impact of Disk Striping on Reliability, *Database Engineering*, Vol. 11, No. 1, 1988, pp. 26-39
- [12] Ghandeharizadeh, S., DeWitt, D. J., A Multiuser Performance Analysis of Alternative Declustering Strategies, *Proceedings of the 6th International Conference on Data Engineering*, 1990, pp. 466-475
- [13] Ghandeharizadeh, S., DeWitt, D. J., Hybrid-Range Partitioning Strategy: A New Declustering Strategy for Multiprocessor Database Machines, *Proceedings of the 16th International Conference on Very Large Data Bases*, 1990, pp. 481-492
- [14] Gibson, G.A., Hellerstein, L., Karp, R.M., Katz, R.H., Patterson, D.A., Failure Correction Techniques for Large Disk Arrays, *Proceedings of the 3rd International Conference on*

Architectural Support for Programming Languages and Operating Systems, 1989, pp. 123-132

- [15] Gifford, D., Spector, A., The TWA Reservation System, *Communications of the ACM*, Vol.27 No.7, 1984
- [16] Graefe, G., Volcano, an Extensible and Parallel Query Evaluation System, *Technical Report No. CU-CS-481-90*, University of Colorado at Boulder, 1990
- [17] Gray, J.N., Horst B., Walker, M., Parity Striping of Disc Arrays: Low-Cost Reliable Storage with Acceptable Throughput, *Proceedings of the 16th International Conference on Very Large Data Bases*, 1990, pp. 148-161
- [18] Haerder, Th., Meyer-Wegner, K., Mitschang, B., Sikeler, A., PRIMA - A DBMS Prototype Supporting Engineering Applications, *Proceedings of the 13th International Conference on Very Large Data Bases*, 1987, pp. 433-442
- [19] R. Heydt and D. Heydt, DB2 Version 2 Database Design and Administration, *John Wiley & Sons*, 1989
- [20] Hornick, M.F., Zdonik, S.B., A Shared, Segmented Memory System for an Object-Oriented Database, *ACM Transactions on Office Information Systems*, Vol. 5, No. 1, 1987, pp. 70-95
- [21] Hsiao, H., DeWitt, D.J., Chained Declustering: A New Availability Strategy for Multiprocessor Database Machines, *Proceedings of the 6th International Conference on Data Engineering*, 1990, pp. 456-465
- [22] Hudson, S. E., King, R., Cactis: A Self-Adaptive, Concurrent Implementation of an Object-Oriented Database Management System, *ACM Transactions on Database Systems*, Vol. 14, No. 3, 1989, pp. 291-321
- [23] Katz, R.H., Ousterhout, J.K., Patterson, D.A., Stonebraker, M.R., A Project on High Performance I/O Subsystems, *Database Engineering*, Vol. 11, No. 1, 1988, pp. 40-47
- [24] Kim, M.Y., Synchronized Disk Interleaving, *IEEE Transactions on Computers*, Vol. C-35, No. 11, 1986, pp. 978-988
- [25] Lehman, T.J., Lindsay, B.G., The Starburst Long Field Manager, *Proceedings of the 15th International Conference on Very Large Data Bases*, 1989, pp. 375-383
- [26] Livny, M., Khoshafian, S., Boral, H., Multi-Disk Management Algorithms, *ACM SIGMETRICS Conference*, 1987, pp. 69-77
- [27] M2344K Micro-Disk Drives CE Manual, *Document No. 41FH6817E-01A*, Fujitsu Ltd., 1987
- [28] Matsliach, G. and Shmueli, O., Maintaining Bounded Disorder Files in Multiprocessor Multi-Disk Environments, *Proceedings of the International Conference on Database Theory*, 1990
- [29] Moad, J., Relief for Slow Storage Systems, *Datamation*, Vol. 36, No. 17, 1990, pp. 20-28

- [30] Muntz, R.R., Lui, J.C.S., Performance Analysis of Disk Arrays Under Failure, *Proceedings of the 16th International Conference on Very Large Data Bases*, 1990, pp. 162-173
- [31] Olson, T.M., Disk Array Performance in a Random IO Environment, *ACM Computer Architecture News*, Vol. 17, No. 5, 1989, pp. 71-77
- [32] Omiecinski, E., Scheuermann, P., A Parallel Algorithm for Record Clustering, to appear in *ACM Transactions on Database Systems*
- [33] Oracle RDBMS Database Administrator's Guide Version 6.0, *Oracle Corp.*, 1988
- [34] Ousterhout, J.K., DaCosta, H., Harrison, D., Kunze, J.A., Kupfer, M., Thompson, J.G., A Trace-Driven Analysis of the UNIX 4.2 BSD File System, *Proc. ACM Symposium on Operating System Principles*, 1985
- [35] Patterson, D.A., Gibson, G., Katz, R.H., A Case for Redundant Arrays of Inexpensive Disks (RAID), *Proceedings of the SIGMOD International Conference on Management of Data*, 1988, pp. 109-116
- [36] Reddy, A.L., Banerjee, P., An Evaluation of Multiple-Disk I/O Systems, *IEEE Transactions on Computers*, Vol. 38, No. 12, 1989, pp. 1680-1690
- [37] Salem, K., Garcia-Molina, H., Disk Striping, *Proceedings of the 2nd International Conference on Data Engineering*, 1986, pp. 336-342
- [38] Samadi, B., TUNEX: A Knowledge-Based System for Performance Tuning of the UNIX Operating System, *IEEE Transactions on Software Engineering*, Vol. 15, No. 7, 1989, pp. 861-874
- [39] Schek, H.-J., Paul, H.-B., Scholl, M.H., Weikum, G., The DASDBS Project: Objectives, Experiences, and Future Prospects, *IEEE Transactions on Knowledge and Data Engineering*, Vol.2 No.1, 1990, pp. 25-43
- [40] Scheuermann, P., Park, Y., Omiecinski, E., Heuristic Reorganization of Clustered Files, *Proceedings of the International Conference on Foundations of Data Organization (FODO)*, 1989, pp. 16-30
- [41] Schwetman, H., CSIM Reference Manual (Revision 13), MCC Technical Report ACA-ST-252-87, Rev. 13, *Microelectronics and Computer Technology Corporation*, Austin, 1989
- [42] Smith, A.J., Input/Output Optimization and Disk Architectures: A Survey, *Performance and Evaluation*, Vol. 1, 1981, pp. 104-117
- [43] Smith, M., Alexander, W., Borat, H., Copeland, G., Keller, T., Schwetman, H., Young, C.-R., An Experiment on Response Time Scalability in Bubba, 6th International Workshop on Database Machines, Springer, 1989
- [44] Stonebraker, M., The Case for Shared Nothing, *IEEE Database Engineering*, Vol. 9, No. 1, 1986, pp. 4-9
- [45] Stonebraker, M., Katz, R., Patterson, D., Ousterhout, J., The Design of XPRS, *Proceedings of the 14th International Conference on Very Large Data Bases*, 1988, pp. 318-330

- [46] Stonebraker, M., Schloss, G.A., Distributed RAID – A New Multiple Copy Algorithm, *Proceedings of the 6th International Conference on Data Engineering*, 1990, pp. 430–437
- [47] The Tandem Database Group, NonStopSQL: A Distributed, High-Performance, High-Availability Implementation of SQL, *2nd Int. Workshop on High Performance Transaction systems*, Springer, Lecture Notes in Computer Science No. 359, 1989, pp. 60–103
- [48] Teradata, DBC/1012 Database Computer System Manual Release 2.0, *Document No. C10-0001-02*, Teradata Corp., 1985
- [49] Weikum, G., Set-Oriented Disk Access to Large Complex Objects, *Proceedings of the 5th International Conference on Data Engineering*, 1989, pp. 426–433
- [50] G. Weikum, C. Hasse, A. Moenkeberg, and P. Zabback, The COMFORT Project: A Comfortable Way to Better Performance, *Technical Report, ETH Zurich*, 1990
- [51] Zabback, P., Paul, H.-B., Deppisch, U., Office Documents on a Database Kernel – Filing, Retrieval, and Archiving, *Proceedings of the International Conference on Office Information Systems*, 1990, pp. 261–270

Appendix

The appendix describes the analytical model that we have developed to determine the optimal degree of parallelism for a read or write request to a particular file. For the sake of simplicity we assume that the file consists of one region. We assume that all disks in the disk array are of the same type. The following abbreviations are used:

- N = number of disks in the disk array
- C = number of cylinders on one disk
- T = number of tracks in one cylinder
- B = number of blocks in one track
- R = average request size of the file (in blocks)
- P = optimal degree of parallelism for serving a single request.

In the following, we will derive a formula for the access time to R consecutive blocks of a file as a function of P . For tractability, we assume single-user mode or contention-free multi-user mode, i.e., disregard queueing delays. We will then derive the degree of parallelism that minimizes the value of this function.

Seek Latency

Suppose that the request is served by P disks. This requires seeks on all P involved disks. The longest of these seeks has to cross $n_{seek}(P)$ cylinders where according to [2]

$$n_{seek}(P) = C \left(1 - \prod_{i=1}^P \frac{2i}{2i+1} \right)$$

The following approximation of $n_{seek}(P)$ with suitable constants $a=0.577$ and $b=-0.118$ leads us to

$$n_{seek}(P) \approx C(1 - a - b \ln(P))$$

Let $t_{seek}(x)$ denote the time to move the disk arm over x cylinders. The function $t_{seek}(x)$ can be approximated by the following linear equation

$$t_{seek}(x) = ex + f$$

with suitable constants e and f and sufficiently large x . Finally, the seek latency for a request is

$$A_{seek}(P) = t_{seek}(n_{seek}(P))$$

Rotational Latency

Now we derive the maximum rotational delay when P disks perform parallel accesses. The rotational delays of P independent disks can be viewed as P independent random variables $X_1, \dots, X_P$ with identical distribution. We assume a uniform distribution in the range $0 \leq X_i \leq 1$ ($i=1, \dots, P$), i.e., the random variables denote the fraction of a full rotation that is necessary to position the head on the first requested block. Now let

$$X_{\max} = \max(X_1, X_2, \dots, X_P)$$

denote the maximum of these random variables, i.e., the maximum rotational delay. For each X_i ($i=1, \dots, P$), we have the probability distribution

$$P(X_i \leq r) = r \quad (0 \leq r \leq 1).$$

From this we obtain

$$P(X_{\max} \leq r) = \prod_{i=1}^P P(X_i \leq r) = r^P$$

and

$$P(X_{\max} = r) = \frac{dP(X_{\max} \leq r)}{dr} = Pr^{P-1}$$

which in turn yields

$$E[X_{\max}] = \int_0^1 r Pr^{P-1} dr = \frac{P}{P+1}.$$

Thus, if t_{rot} is the time for one disk revolution, the rotational latency of P disks is

$$A_{rot}(P) = t_{rot} \frac{P}{P+1}.$$

For $P = 1$ (i.e., only one disk is involved), we obtain the usual result $A_{rot}(1) = \frac{t_{rot}}{2}$.

Data Transfer Time

The P disks that serve a request of size R transfer data in parallel. Let $d = \left\lceil \frac{R}{P} \right\rceil$ be the size of the portion that is transferred by one disk. If these d blocks span track boundaries, then the necessary head switching involves head fine-positioning on high-density disks.

This fine-positioning causes an additional delay since it requires reading servo information that is embedded in the tracks. If consecutive tracks are properly offset, this head switching delay is typically a few milliseconds. Let t_{hs} denote this delay. For $d \leq B$, the probability that d blocks span more than one track is

$$1 - \frac{B-d+1}{B} = \frac{d-1}{B}$$

Thus, we can estimate the number of necessary head switches n_{hs} by (omitting the ceiling function and rather using the continuous approximation)

$$n_{hs} = \left\lceil \frac{d}{B} \right\rceil - 1 + \frac{d - \left(\left\lceil \frac{d}{B} \right\rceil - 1\right)B - 1}{B} \approx \frac{R}{PB} - 1 + \frac{\frac{R}{P} - \left(\frac{R}{PB}\right)B + B - 1}{B} = \frac{R-P}{PB}.$$

Some of these head switches occur in combination with a cylinder switch. Let t_{cs} denote the time for this cylinder switch (t_{cs} is the time for a seek of length 1). Analogously to the number of head switches we can estimate the number of cylinder switches n_{cs} by

$$n_{cs} = \left\lceil \frac{n_{hs} + 1}{T} \right\rceil - 1 + \frac{T-H}{TB} \approx \frac{2\frac{R}{P} - 2 + B - TB}{TB}$$

Then, the transfer time for a request of size R with parallelism P is

$$\begin{aligned} A_{transfer}(P) &= (n_{hs} - n_{cs})t_{hs} + n_{cs}t_{cs} + \frac{R}{PB}t_{rot} \\ &= \left(\frac{R-P}{PB} - \frac{2\frac{R}{P} - 2 + B - TB}{TB} \right) t_{hs} + \frac{2\frac{R}{P} - 2 + B - TB}{TB} t_{cs} + \frac{R}{PB} t_{rot}. \end{aligned}$$

Total Access Time

Combining the above results yields the following total access time:

$$\begin{aligned} A_{total}(P) &= A_{seek}(P) + A_{rot}(P) + A_{transfer}(P) \\ A_{total}(P) &= eC(1-a-b\ln(P)) + f + \frac{Pt_{rot}}{P+1} + \frac{R-P}{PB}t_{hs} + \frac{2\frac{R}{P} - 2 + B - TB}{TB}(t_{cs} - t_{hs}) + \frac{R}{PB}t_{rot} \end{aligned} \quad (1)$$

For the degree of parallelism P_{opt} that minimizes $A_{total}(P)$, the following equation must hold:

$$\frac{dA}{dP}(P_{opt}) = \frac{t_{rot}}{P_{opt} + 1} - \frac{eCb}{P_{opt}} - \frac{P_{opt}t_{rot}}{(P_{opt} + 1)^2} - \frac{t_{hs}}{P_{opt}}B - \frac{(R-P)t_{hs} + Rt_{rot}}{P_{opt}^2B} - 2\frac{R(t_{cs} - t_{hs})}{P_{opt}^3TB} = 0$$

This equation is equivalent to

$$\begin{aligned} &-eCbBTP_{opt}^3 + (t_{rot}TB - RT(t_{hs} + t_{rot}) - 2eCbBT + 2R(t_{hs} - t_{cs}))P_{opt}^2 \\ &+ (-2RT(t_{rot} + t_{hs}) - eCbBT + 4R(t_{hs} - t_{cs}))P_{opt} + R(2(t_{hs} - t_{cs}) - T(t_{rot} + t_{hs})) = 0 \end{aligned} \quad (2)$$

Solving this cubic equation by using Cardan's formula finally leads us to a solution for the optimal degree of parallelism. For the sake of simplicity, this is shown only for a specific disk.

Solution for a specific disk

For the Fujitsu M2344K drive [27], which is representative for most modern disk drives we have $e = 0.03435$, $C = 624$ cylinders, $T = 27$ tracks, $t_{hs} = 0.2$ ms, $t_{cs} = 5$ ms, $t_{rot} = 16.66$ ms, and $B = 10$ blocks (with block size 4 KB). Thus, formula (2) becomes

$$c_1 P_{opt}^3 + c_2 P_{opt}^2 + c_3 P_{opt} + c_4 = 0 \quad (3)$$

with $c_1 = 2731.6$, $c_2 = -465.25R + 23463.2$, $c_3 = -930.5R + c_1$, $c_4 = -465.25R$. Figure 13 shows a plot of this function for the request size $R = 25$ blocks.

Figure 13: Plot of function (3)

The equation (3) has three real solutions, but only one is positive. Also, computing points close to this positive solution confirms that we found the minimum of $A_{total}(P)$. Thus our solution is:

$$P_{opt}(R) = c_5 R + 2 \sqrt{-\frac{1}{27}(c_6 R^2 + c_7 R + c_8)^3} \cos \left(\frac{\pi}{3} - \frac{1}{3} \arccos \left(\frac{c_9 R^3 + c_{10} R^2 + c_{11} R + c_{12}}{2 \sqrt{-\frac{1}{27}(c_6 R^2 + c_7 R + c_8)^3}} \right) \right) + c_{13} \quad (4)$$

with $c_5 = 0.056$, $c_6 = -0.009$, $c_7 = 0.635$, $c_8 = -23.593$, $c_9 = -0.0004$, $c_{10} = 0.036$, $c_{11} = -1.93$, $c_{12} = 44.08$, and $c_{13} = -2.863$. Figure 14 shows the function $P_{opt}(R)$ for different track sizes B (in blocks).

Figure 14: Optimal degree of parallelism $P_{opt}(R)$ for a request of size R

Throughput Considerations

Let N be the number of disks in the disk array, $\bar{R}$ be the average request size for all files, and $\bar{P}$ be the average degree of parallelism. An average request then has the following device-busy time

$$\begin{aligned} t_{busy} &= \bar{P} \left(t_{seek} \left(\frac{C}{3} \right) + \frac{t_{rot}}{2} + A_{transfer}(\bar{P}) \right) \\ &= \bar{P} \left(t_{seek} \left(\frac{C}{3} \right) + \frac{t_{rot}}{2} + \frac{\bar{R} - \bar{P}}{\bar{P}B} t_{hs} + \frac{2\frac{R}{\bar{P}} - 2 + B - TB}{TB} (t_{cs} - t_{hs}) + \frac{\bar{R}}{\bar{P}B} t_{rot} \right) \end{aligned}$$

Note that the *average* seek and rotational latency of one disk is independent of $\bar{P}$ (as opposed to the *maximum* latency of the $\bar{P}$ disks that serve a request, see above). From τ we can derive the maximum throughput, which is obtained for $\bar{P} = 1$. The throughput τ is calculated as

$$\begin{aligned} \tau &= \frac{N}{t_{busy}} = \frac{N}{\bar{P} \left(t_{seek} \left(\frac{C}{3} \right) + \frac{t_{rot}}{2} + A_{transfer}(\bar{P}) \right)} \\ \tau_{max} &= \frac{N}{t_{seek} \left(\frac{C}{3} \right) + \frac{t_{rot}}{2} + A_{transfer}(1)} \end{aligned}$$

A given throughput requirement τ^0 imposes the constraint $\tau_{max} \geq \tau^0$, which is equivalent to

$$\bar{P} \leq$$

$$\begin{aligned} &\frac{(4\bar{R}(1-k) - 4 + 2B(1-T))t_{cs} + (2T(\bar{R} - 1 + B - k\bar{R}) + 4\bar{R}(k-1) + 4 - 2B)t_{hs} + 2TBt_{seek} + (B + 2\bar{R}(1-k))Tt_{rot}}{k((4 - 2T - 2B + 2TB)t_{hs} - (4 + 2TB - 2B)t_{cs}) + 2TBt_{seek} + TBt_{rot}} \\ &=: \bar{P}_{max} \end{aligned}$$

with $k := \tau^0 / \tau_{max}$. This formula states that, in order to achieve the required throughput τ^0 , the average degree of parallelism must not be higher than $\bar{P}_{max}$. For example, for the Fujitsu M2344K drive, $N = 12$, $\bar{R} = 600\text{KB}$, we get $\tau_{max} \approx 38$ requests/second. For a required throughput of $\tau^0 = 30$ requests/second, we obtain $\bar{P}_{max} = 4$.

Gelbe Berichte des Departements Informatik

129	L. Pfau	GFT: A Tool for Data Management in the UNIX Environment
130	S. Bondeli	Divide and Conquer: A Parallel Algorithm for the Solution of a Tridiagonal Linear System
131	H. Goorhuis, N. Gürman, M. Montigel, L. Thalmann	Neuronale Netze und regelbasierte Systeme: Ein hybrider Ansatz
132	B. Jiang	Design, Analysis, and Evaluation of Algorithms for Computing Partial Transitive Closures in Databases
133	J. Tempel	SPARC-Oberon User's Guide
134	G. Wong	An Approach on How to Combine Object Recognition Methods
135	C. Wieland	Two Explanation Facilities for the Deductive Database Management System DeDEx
136	H.-J. Schek, M.H. Scholl, G. Weikum	From the KERNEL to the COSMOS: The Database Research Group at ETH Zürich
137	G. Weikum, C. Hasse, A. Mönkeberg, P. Zaback	The COMFORT Project: A Comfortable Way to Better Performance
138	J. Gutknecht	The Oberon Guide: System Release 1.2
139	P.E. Saylor, D.E. Smolarski	Implementation of an Adaptive Algorithm for Richardson's Method
140	N. Wirth	Die Programmiersprache Oberon
141	M. Franz	The Implementation of MacOberon
142	M. Franz	MacOberon Reference Manual
143	N. Wirth	From Modula to Oberon
The Programming Language Oberon
(revised edition)
144	J.L. Marais	The GADGETS User Interface Management System
145	J. Mössenböck	She: A Simple Hypertext Editor for Programs
146	H.E. Meier	Schriftgestaltung mit Hilfe des Computers
Typographische Grundregeln

Eidg. Techn. Hochschule Zürich

Informatikbibliothek

ETH-Zentrum

CH-8092 Zürich