

Priority-Hints: An Algorithm for Priority-Based Buffer Management

Rajiv Jauhari, Michael J. Carey, and Miron Livny

Computer Sciences Department
University of Wisconsin
Madison, WI 53706

ABSTRACT — In this paper, we address the problem of buffer management in a DBMS when the workload consists of transactions of different priority levels. We present Priority-Hints, a new buffer management algorithm that uses hints provided by the DBMS access methods. The performance of Priority-Hints is compared to that of priority buffer management schemes introduced earlier for a variety of workloads. Our simulation results indicate that Priority-Hints performs consistently better than simple LRU-based algorithms. Furthermore, our algorithm approaches (and in some cases surpasses) the performance of highly sophisticated algorithms that require much more information to be provided to the buffer manager.

1. INTRODUCTION

1.1. Motivation

Priority scheduling has recently become an area of increased interest to the database community [SIGM88, Abbo88, Abbo89, Care89, Hari90]. Applications that require different levels of system response for different transactions (for example, a system that is designed to provide faster service to interactive jobs than to batch jobs) can benefit from priority scheduling at the DBMS resources, as shown in [Care89]. Several data-intensive applications such as computer-aided manufacturing, stock trading, and command and control systems may require real-time response, which can also be supported with the help of priority scheduling at the resources of the DBMS [SIGM88, Abbo89].

The use of priority in DBMS resource scheduling may lead to an increase in the extent to which buffer management impacts system performance compared to its impact in conventional database systems. Unpredictable bursty arrivals of high-priority transactions may force a priority-oriented DBMS to operate in

This research was supported in part by the National Science Foundation under grant IRI-8657323 and by the Digital Equipment Corporation through its Initiatives for Excellence program.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.

Proceedings of the 16th VLDB Conference
Brisbane, Australia 1990

regions where the total load on the buffer pool (i.e., the sum of the buffering requirements of transactions of all priority levels) exceeds the buffer pool capacity. In these operating regions, priority-based load control and buffer allocation policies will be required, as the use of conventional load control and allocation techniques may lead to situations of "priority-inversion," where high-priority transactions are forced to wait while low-priority transactions are allowed to make progress. Furthermore, the set of concurrently active transactions in these operating regions may include transactions of different priority levels. In this scenario, then, priority-based buffer replacement policies may also be required in order to provide preferential service to high-priority transactions. We anticipate that all aspects of buffer management (load control, allocation, and replacement) will become both more complex and more significant when priority is used in scheduling DBMS resources.¹

Several interesting new issues arise when buffer management decisions have to include priority considerations. One such issue is the tradeoff between the overheads introduced as a consequence of the use of priority and the advantages provided to high-priority transactions. If a buffer containing data accessed by a transaction is replaced as a consequence of priority, its data may have to be re-read from disk once the transaction resumes execution. The total load on the system may increase purely as a consequence of the use of priority in buffer management, and alternative priority-based buffer replacement and allocation policies may result in different relative increases in system load.

A second issue of interest is the extent to which information about the workload can be used by the buffer manager to improve system performance in the presence of priority. Existing buffer management schemes assume different levels of information about transactions' data access patterns [Effe84, Teng84, Chou85, Sacc86]. In this paper, we introduce a new buffer management algorithm that makes use of hints provided by the database access methods (as in the Starburst buffer manager [Haas90] and DB2 [Teng84]). This new algorithm, called "Priority-Hints," uses these hints to make priority-based buffer management decisions while trying to minimize the priority-induced overhead on the system.

A third issue in priority-based buffer management is inter-transaction buffering interference across priority levels. For

¹This is in contrast to the trends in conventional database systems, where it may be argued that by increasing the buffer pool size with respect to the database size, and by keeping the multiprogramming level under a certain threshold, buffer management policies can be made essentially irrelevant.

example, update-intensive transactions may quickly make large numbers of buffers "dirty," making them unavailable for replacement until they are written out to disk. The performance of high-priority transactions can thus be affected adversely by low-priority updates. Another example of inter-transaction effects in the presence of priority may occur when high-priority sequential scans quickly replace a large number of buffers with pages that are accessed just once, unnecessarily depriving lower priority non-sequential transactions of buffers that need to be accessed repeatedly. Priority-based buffer management policies should be designed to minimize these effects.

Finally, the importance of using priority-based buffer management in a DBMS that already uses priority at the CPUs and the disks may itself be open to question. In this paper, we investigate all of these issues.

1.2. Related Work

A number of buffer management strategies for conventional database systems have been proposed in the literature [Effe84, Chou85, Sacc86]. Simple techniques such as Global-LRU assume no knowledge of the data access patterns, while algorithms such as Hot Set [Sacc86] and DBMIN [Chou85] attempt to take advantage of the limited number of ways in which queries access data in relational database systems. In [Chou85], it was shown that DBMIN, a relatively sophisticated buffer management strategy in which data access pattern information is supplied to the buffer manager on a per-file basis by the DBMS optimizer, performs better than most other existing buffer management strategies.

In [Teng84], the design of IBM's DB2 buffer manager is described, and several techniques used to maximize buffer pool performance are discussed. For example, DB2's buffer manager distinguishes between sequential accesses and random accesses, and buffered pages that are part of sequential accesses are chosen as replacement victims in preference to randomly accessed pages. The LRU replacement policy is used within each group of buffers (sequential or random). Other performance-enhancing techniques used in DB2 include the use of prefetch for sequential scans and the use of deferred-write mechanisms for writing updated pages to disk. In the Starburst system, as described briefly in [Haas90], access methods can provide hints to the buffer manager about their expected future re-use of each buffered page in order to guide replacement decisions.

An initial investigation of the problem of priority scheduling at the physical resources of a DBMS was described in [Care89]. Algorithms for priority-based disk scheduling and CPU scheduling were presented, and priority-oriented modifications of two existing buffer management algorithms (Global-LRU and DBMIN) were described. It was shown that a single priority scheduler does not suffice to meet the goals of priority scheduling in a DBMS due to the heterogeneity and the multiplicity of its resources. Priority scheduling at the admission control component of the DBMS (the part of the system that determines whether an arriving transaction is to be allowed to enter the system, or is to be blocked outside the system until sufficient buffer resources become available) was shown to be particularly

important in order to keep the system stable from the point of view of high-priority transactions as system load is increased. A homogeneous workload model was used to compare the performance of the priority-scheduling algorithms, and Priority-DBMIN (the priority version of DBMIN) was shown to perform as well as or better than Priority-LRU (the priority version of Global LRU).

In [Abbo89], priority scheduling at the CPU and the disk, as well as for concurrency control, was shown to be effective in reducing the number of transactions that miss their deadlines in a real-time DBMS. Transactions were modeled as random sequences of page accesses, and the buffer pool was modeled simply by computing the probability of finding a page in main memory assuming a uniform probability of access to the entire database. Finally, it was shown in [Hari90] that in a real-time environment, priority scheduling at the CPU and the disks may help optimistic concurrency control algorithms outperform locking schemes. Issues of buffer replacement strategies and workloads with non-random data access patterns were not included in either [Abbo89] or [Hari90].

1.3. Our Work

The design of the Priority-Hints algorithm was motivated by the intuition that the use of simple page-level information by the buffer manager (as in [Teng84, Haas90]) may improve system performance over that provided by simple LRU-based approaches. Such a performance improvement acquires great importance in priority-based systems, where significant I/O overheads may result as a consequence of using priority in buffer management decisions. Improvements in performance (relative to simple LRU-based approaches) may also be obtained by using sophisticated DBMIN-like algorithms, as shown in [Care89], but only at the cost of added system complexity; our goal is to examine whether such complexity is really required.

In order to take a detailed look at the issues involved in priority-based buffer management in the presence of mixed workloads, we have implemented a simulation model of a DBMS that uses priority scheduling. In this paper, we use this simulation model to compare the performance of our new algorithm with Priority-LRU and Priority-DBMIN. Our primary objective in this analysis is to explore the extent to which our algorithm can surpass the performance of Priority-LRU, and how close we can get to the performance of Priority-DBMIN. We also conduct experiments that shed light on the other priority-related buffer management issues raised earlier.

The remainder of this paper is organized as follows: In Section 2, we describe Priority-Hints, our new buffer management algorithm, and review the Priority-LRU and Priority-DBMIN policies introduced in [Care89]. Section 3 describes our simulation model of a priority-oriented DBMS. Section 4 presents a series of performance experiments that help us to understand the tradeoffs involved in using priority in buffer management. Finally, Section 5 summarizes our contributions and describes our plans for future work.

2. PRIORITY-BASED BUFFER MANAGEMENT ALGORITHMS

In this section we present the Priority-Hints algorithm. Our assumptions about buffer management are outlined first. We then describe Priority-Hints and briefly review Priority-LRU and Priority-DBMIN, the two algorithms presented in [Care89]. Our scheme for handling dirty data, which is common to the three algorithms, is described next; and we conclude the section with a summary of the key differences between the three algorithms.

2.1. Buffer Management Assumptions

A page is assumed to be fixed (or pinned) in the buffer pool during the interval when a transaction is processing the data on the page. As soon as the transaction has finished processing the page, it unfixes it. Fixed pages cannot be chosen as buffer replacement victims. The *owner* of a resident page is the transaction with the highest priority among the executing transactions that have accessed the page since it was brought into memory. The buffer manager associates a *timestamp* with each resident page in order to keep track of the recency of usage of pages. Each time the data in a buffer is accessed or updated, a global counter is incremented and its new value is inserted as the timestamp of the page. Thus, the larger the value of the timestamp of a page, the more recently the page was accessed. Pages in the free list (and the dirty list, described at the end of this section) are kept in LRU order using their timestamps.

Based on the number of buffers available, a transaction may be admitted to the system right away, or it may be blocked initially. Transactions blocked outside the system are queued in order of priority. Once a transaction is allowed to begin execution, it continues until it commits, is aborted as a result of concurrency control, or is *suspended*.² A transaction is said to be *suspended* by the buffer manager if it is temporarily prohibited from making further buffer requests; the buffers owned by the transaction are freed. The buffer manager considers *reactivating* suspended transactions at the same decision point that it considers admitting blocked transactions, which is whenever a running transaction completes or aborts. A reactivated transaction resumes its execution at the point where it was suspended.

A transaction checks whether it has been chosen as a suspension victim at instants when it has no pages fixed. We call such instants "suspension-safe" points. At suspension-safe points, the transaction "volunteers" to let all of its buffers be stolen by transactions of higher priority. Such instants occur normally during the execution of a transaction. In a sequential scan, for example, they occur at the point when the transaction unfixes one page and is about to request that another be fixed. In addition, a priority DBMS should be designed so that transactions periodically "come up for air" and check if they need to give up their buffers.

²A running transaction may also be blocked temporarily if there are no free buffers available, and all of the in-use buffers are either pinned or dirty. The transaction's buffer request is then enqueued in a queue called the *Buffer Waiting Queue*. Queued buffer requests are served in priority order.

2.2. The Priority-Hints Algorithm

As its name suggests, Priority-Hints makes use of hints (provided by the DBMS access methods) that indicate whether a particular data page should be retained in memory in preference to other data pages. The basic ideas underlying Priority-Hints are the following:

- As discussed in [Teng84, Haas90], it is possible to classify the pages referenced by a transaction into two groups: pages that are likely to be re-referenced by the same transaction (such as the pages of the inner file in a nested-loops join), and pages that are likely to be referenced just once (such as the pages of a file being scanned sequentially). The pages that are likely to be re-referenced are called *favored* pages, and the others are called *normal* pages. We assume that whenever a request for a page is made to the buffer manager, the buffer manager is informed whether the requested page is favored or normal.
- The favored pages of a transaction should be kept in the buffer pool as long as the transaction needs to reaccess them; each normal page should be made available for replacement as soon as the transaction unfixes it. When searching for replacement victims, normal pages should therefore be considered before favored pages.
- If it becomes necessary to choose a favored page as a replacement victim, the most-recently-used (MRU) policy should be used to choose the victim. As discussed in [Chou85], MRU is a better approach than LRU when choosing replacement victims from a set of pages that are being repeatedly looped over, and favored pages are likely to fall into this category.

The Priority-Hints algorithm combines these ideas with the notion of priority as follows.

Buffer Pool Organization: Buffers are organized into "transaction sets," where a transaction set consists of all of the buffers owned by a single transaction.³ Transaction sets are arranged in priority order, with recency of arrival of the owner transaction being used to break ties if there are multiple transactions of the same priority. In the buffer pool configuration shown in Figure 2.1, there are three transactions (T1, T3, and T2), three priority levels, and no free buffers.

A transaction set consists of two kinds of buffers: the buffers currently fixed by the owner (marked by the letter "F" in Figure 2.1), and buffers containing unfixed favored pages of the owner (marked by the letter "U" in Figure 2.1). The unfixed favored pages are maintained in MRU order with the help of buffer timestamps. Note that a transaction set contains no unfixed normal pages; whenever a normal page is unfixed, it is freed.

Transaction Admission: Transactions are required to estimate the maximum number of pages that they will need to fix concurrently, and the buffer manager keeps track of the sum of these "fixing requirements" for all active transactions. If admitting a newly arrived transaction does not cause this sum to exceed the

³Buffers containing pages shared by more than one transaction are owned by the transaction with the highest priority among the sharers.

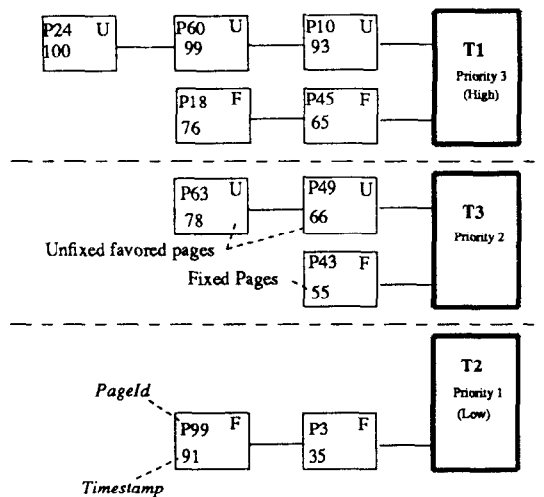


Figure 2.1: Example of Priority-Hints Buffer Pool Organization.

size of the buffer pool, the transaction is admitted. Otherwise, if there are running transactions of lower priority than that of the new arrival, the one(s) with the lowest priority among them are suspended until there are enough buffers for the new arrival, or until no lower priority transactions remain.⁴ If no remaining transactions are of a priority less than the new arrival, then the new arrival is forced to wait outside the DBMS.

Buffer Replacement and Allocation: When a buffer miss occurs and there is no free page available, the buffer manager first attempts to get a replacement victim from among the unfixed favored pages of transactions of lower priority than the requesting transaction. The buffer pool searches its transaction sets in inverse priority order, starting from the lowest priority transaction, looking for unfixed favored pages. It stops searching on either of the following conditions:

- (1) It finds a transaction of lower priority than the requesting transaction with an unfixed favored page; or
- (2) it has reached a transaction of a priority equal to or greater than that of the requesting transaction.

In case (1), it chooses the most recently unfixed favored page of the lower-priority transaction as the replacement victim. In case (2), it chooses the most recently unfixed favored page (if any) of the *requesting transaction* itself. Note that this means that transactions cannot steal buffers from other transactions of the same priority; thus the replacement policy for favored pages is *local* rather than *global*. If no replacement victim is available, then the outstanding request is queued in the Buffer Waiting Queue. Furthermore, if there are running transactions of lower priority, the transaction with the lowest priority among them is

⁴In choosing suspension victims from among transactions of the same priority, later arrivals are chosen for suspension in preference to earlier arrivals. Also, earliest arrival time is the criterion for choosing the transaction (from among a group of waiting transactions of the same priority) that should first attempt to enter the system.

suspended. Continuing the example of Figure 2.1, if T1 makes a buffer request for page P6, which is not in the buffer pool, the buffer manager will start its search for replacement at T2. Finding no unfixed buffer in T2's transaction set, it will look at T3's transaction set and find P63 as the replacement victim. Had there been no unfixed pages of priority 1 or 2, then P100 would have been chosen as the replacement victim.

To summarize, Priority-Hints has a *local MRU* replacement policy for favored pages, and a *global LRU* replacement policy for normal pages. (Recall that normal pages are placed on the free list at unfix time, and that the free list is maintained in LRU order.)

2.3. The Priority-LRU Algorithm

In Priority-LRU, the prioritized version of Global LRU, the buffer pool is organized dynamically into priority levels as described in [Care89]. Each priority level consists of pages whose owners have the same priority, and the pages within a level are arranged in LRU order. The transaction admission policy for Priority-LRU is the same as that for Priority-Hints. The key idea of the Priority-LRU replacement policy is that the least recently unfixed page of the lowest priority should be chosen as the victim. If there are no free buffers, the search for a replacement victim starts at the lowest priority queue, where we check whether unfixed candidates are available. If such candidates are found, the least-recently-used candidate is chosen as the victim. If no candidate is found at this priority level, we move up one level, and we repeat the process until we have either found a victim, reached a priority level that exceeds that of the requesting transaction, or exhausted the search. If no victim is found, and there are transactions of lower priority running, the lowest-priority transaction is suspended as in Priority-Hints.

2.4. The Priority-DBMIN Algorithm

As discussed in [Chou85], the primitive operations (e.g., selections, joins) of transactions in a relational DBMS can be described as a composition of a set of regular reference patterns such as sequential scans and hierarchical index lookups. These patterns are known to the query optimizer. The DBMIN buffer management policy makes use of this information in the following way: A set of buffers (called a "locality set") is allocated to each transaction for each file accessed by it. The optimum size of each locality set and the optimum replacement policy to be used within a locality set are supplied to the buffer manager by the optimizer. DBMIN guarantees that each transaction that is allowed to enter the system has the optimum number of buffers available to it, and the optimum replacement policy is used within each locality set.

Priority-DBMIN, the prioritized version of DBMIN, also allocates buffers to transactions in locality sets. A transaction is allowed to enter the system only if its optimal-sized locality sets can be accommodated in the buffer pool. Otherwise, if there are transactions of lower priority than that of the arriving transaction in the system, they are suspended in reverse priority order until sufficient buffers become available for the new arrival. As in the original DBMIN algorithm, Priority-DBMIN uses the optimizer-

supplied optimum replacement policy within each locality set.

2.5. Dirty Data

In all three algorithms, a process called the *asynchronous write engine* [Teng84] is responsible for flushing dirty buffers to disk. When a transaction frees a buffer, the buffer is inserted into the free list if it is clean (i.e., if it has no update that has not been written to disk). If the data in the buffer has been updated, the buffer is placed in a queue called the dirty list. The write engine is activated periodically; it also wakes up whenever a buffer miss occurs and the free list is empty. The engine flushes each page in the dirty list that is sufficiently "old" in terms of its recency of use.⁵

Requests to write dirty buffers are asynchronous. This may result in some buffer requests having to wait until a buffer is flushed to disk. Write requests to the disk are therefore assigned a priority equal to the highest possible transaction priority. When its I/O is completed, a dirty buffer is marked clean and placed on the free list, and if there are any buffer requests pending, the highest-priority request is serviced. When choosing replacement victims, dirty data is avoided as long as possible. That is, if a buffer that would normally be a candidate for replacement is dirty, we ignore it in our search for replacement victims unless all candidate buffers are dirty.⁶

2.6. Discussion

In summary, the key features of Priority-Hints that distinguish it from the other algorithms discussed are the following:

- By realizing which pages are normal, Priority-Hints is able to free more buffers earlier in the course of a transaction's execution than Priority-LRU. In this respect, Priority-Hints behaves similarly to Priority-DBMIN.
- When choosing replacement victims from among non-free pages, Priority-LRU chooses the least-recently-unfixed page; Priority-Hints chooses the most-recently-unfixed page. MRU is likely to be a better policy when the replacement victim is part of a set of pages that are being looped over; in cases where pages are reaccessed randomly, the performance differences between MRU and any other replacement policy are negligible [Chou85]. Note that Priority-Hints uses MRU only for favored pages, where it may be advantageous to do so; LRU is still used for normal pages, since normal pages are placed in LRU order in the free list as soon as they are unfixed.
- Priority-Hints' replacement policy ensures that the favored pages of a transaction can be stolen only by a transaction of higher priority: in Priority-LRU, transactions of the same priority

⁵Timestamp information is used to determine whether a dirty page is "old" enough; this check is performed to prevent unnecessary writes of pages that are frequently updated.

⁶Potential deadlocks caused by the entire set of possible replacement candidates being dirty are avoided by synchronously writing dirty buffers to disk in this exceptional situation.

can steal each other's pages.⁷ In Priority-DBMIN, in contrast, a transaction protects its favored pages throughout its execution; if it is not possible to protect them, the transaction is suspended. Note that Priority-Hints allows a transaction to execute even when it does not have an optimum number of favored pages, while Priority-DBMIN does not.

- Priority-LRU does not discriminate between transactions of the same priority when choosing replacement victims. The local replacement search strategy of Priority-Hints, however, ensures that among transactions of the same priority, all but the latest arrival will execute undisturbed as long as the latest arrival has some unfixed favored buffers. Thus, the performance degradation caused by stealing favored buffers is limited to one transaction at a time in Priority-Hints.

- Priority-Hints does not require that information such as optimum locality set sizes be provided by the optimizer, as Priority-DBMIN does. It merely requires that hints be provided to distinguish between normal and favored pages; similar hints are provided in existing DBMSs such as DB2 [Teng84] and Starburst [Haas90]. Thus, Priority-Hints requires less information than Priority-DBMIN.

- Finally, note that while the information supplied to the buffer manager is similar in Priority-Hints, DB2, and Starburst, Priority-Hints differs from the DB2 and Starburst buffer management algorithms in two significant respects. Firstly, it groups buffers on a per-transaction basis in order to allow a local replacement search strategy. Secondly, unfixed buffers are arranged in MRU order in Priority-Hints, unlike in DB2 or Starburst. These two factors may have a significant impact on performance, as will become clear in Section 4.

3. MODELING A PRIORITY-ORIENTED DBMS

In this section, we describe our performance model of a priority-oriented DBMS. The model, which we implemented using the DeNet simulation language [Livn89], consists of five components: the database itself; a *Source*, which generates the workload of the system; a *Transaction Manager*, which models the execution behavior of transactions; a *Resource Manager*, which models the CPU, I/O, and buffer resources of the system; and a *Concurrency Control Manager*, which implements the details of a particular concurrency control algorithm. Since we will be using workloads where concurrency control is not an issue, we will not discuss the Concurrency Control Manager further. (As described in Section 4, our workloads consist either of read-only transactions with data sharing, or updates without data sharing.) In most respects, our model is similar to the model described in [Care89]. Therefore, we describe its components very briefly here; see [Care89] for more details.

3.1. Modeling the Database

The database is modeled as a collection of *relations*. Each relation in turn is modeled as a collection of pages. Indices (clustered or unclustered B+ Trees) on the base relations are

⁷It is advisable to allow Priority-LRU to do this, as many of the pages owned by a transaction are likely to be accessed just once.

included in the database model. The parameters for the database model are summarized in Table 3.1.

3.2. The Source Module

The Source module is the component responsible for modeling the workload of the DBMS. Table 3.2 summarizes the key parameters of the workload model. A transaction may belong to any one of *NumClasses* classes, and it may have any one of *NumPriorities* priority levels. The model is that of an open system, and transactions of each $\langle class_i, priority_j \rangle$ combination arrive at the system in a Poisson process with a mean arrival rate of $ArrRate_{ij}$. Transactions can be single-relation selects, single-relation select-updates, or two-relation select-joins; the type of a transaction of class *i* and priority *j* is controlled by $TransType_{ij}$. Selections can be performed via sequential scans or index scans, and we model three join methods: nested loops joins, classic hash joins, and index joins.

For each transaction type (selection, join, or update) of a particular priority level, an execution plan is provided in the form of a set of parameters. For selections, the access path and the mean selectivity are provided as parameters. The actual selectivity is varied uniformly over the range $[Selectivity_{ijk}/2, 3*Selectivity_{ijk}/2]$. For select-joins, the join method and the inner and outer relations are provided in addition to the selection parameters. For select-updates, the probability of updating a page is specified as the parameter $UpdateProb_{ijk}$. Finally, times spent at the CPU for processing or updating a page are uniformly distributed: the CPU time per data page of relation *k* varies uniformly over the range $[DataPageCPU_{ijk}/2,$

<i>NumRelations</i>	Number of relations in database
<i>RelSize_i</i>	Number of pages in relation <i>i</i>
<i>Indexed_i</i>	Whether relation <i>i</i> has an index
<i>IndexType_i</i>	Type of index (clustered/nonclustered)
<i>Fanout_i</i>	Fanout of internal nodes of index

Table 3.1: Database Model Parameters.

Overall Arrival Pattern Parameters	
<i>NumClasses</i>	Number of transaction classes
<i>NumPriorities</i>	Number of transaction priority levels
Per (Class, Priority) Parameters ($1 \leq i \leq NumClasses, 1 \leq j \leq NumPriorities$)	
<i>ArrRate_{ij}</i>	Transaction Arrival Rate
<i>TransType_{ij}</i>	Transaction Type
<i>JoinMethod_{ij}</i>	Join algorithm used
<i>Outer_{ij}</i>	Outer relation
<i>Inner_{ij}</i>	Inner relation
<i>AccessPath_{ijk}</i>	Access path used to access <i>k</i> th relation
<i>Selectivity_{ijk}</i>	Fraction of <i>k</i> th relation selected
<i>UpdateProb_{ijk}</i>	Probability of page update
<i>IndexPageCPU_{ijk}</i>	CPU time for processing an index page
<i>DataPageCPU_{ijk}</i>	CPU time for processing a data page
<i>UpdateCPU_{ijk}</i>	CPU time for page update

Table 3.2: Workload Model Parameters.

$3*DataPageCPU_{ijk}/2]$, and similar distributions are used for $IndexPageCPU_{ijk}$ and $UpdateCPU_{ijk}$. Given a plan, the Source module generates a list of page accesses that models the sequence in which pages will be accessed by the transaction.

3.3. The Transaction Manager Module

The Transaction Manager is responsible for accepting transactions from the Source and modeling their execution. For each page accessed by the transaction, the Transaction Manager sends a read (or write) request to the Resource Manager; the Resource Manager informs the Transaction Manager when the request is completed. The Resource Manager also informs the Transaction Manager when a transaction is suspended or reactivated. When the Resource Manager decides to reactivate a suspended transaction, the Transaction Manager ensures that the reactivated transaction resumes execution at the point where it was suspended.

3.4. The Resource Manager Module

The Resource Manager controls the physical resources of the DBMS, including the CPU, the disk, and the buffer pool in main memory. Three versions of the Resource Manager have been implemented, supporting the Priority-LRU, Priority-DBMIN, and Priority-Hints algorithms, respectively.⁸ Resource Manager parameters are summarized in Table 3.3.

CPU and Disk Models: The DBMS has *NumCPUs* CPUs and a single priority queue for outstanding CPU requests. The actual CPU where a request is processed is selected at random from among the idle CPUs, if any. The length of each CPU request from a transaction is its per-page CPU processing time; each transaction voluntarily gives up the CPU after processing or updating one page, as in the priority-based round robin CPU scheduling scheme described in [Care89]. There are *NumDisks* disks in the system, with requests at each disk being priority-scheduled according to the prioritized elevator algorithm [Care89]. We model the data as being uniformly distributed across all disks and across all tracks within a disk. The total time required to complete a disk access is computed as the sum of its seek time, rotational latency, and transfer time components. As in [Bitt88, Care89], there is a square root relationship relating seek time to seek distance, and the rotational latency and transfer time are together modeled as a single parameter called *DiskConst*.

Buffer Manager Models: The Buffer Manager component of the resource manager encapsulates the details of the buffer management scheme employed. The number of page frames in the buffer pool is specified as *NumBuffers*. A separate buffer manager component has been implemented for each buffer management algorithm studied.

⁸In addition, Resource Managers supporting Global-LRU and Global-Hints (an algorithm introduced in Section 4) have also been implemented.

<i>NumCPUs</i>	Number of CPUs
<i>NumDisks</i>	Number of disks
<i>NumTracks</i>	Number of tracks per disk
<i>DiskConst</i>	Sum of rotational and transfer delays
<i>SeekFactor</i>	Factor relating seek time to seek distance
<i>NumBuffers</i>	Number of buffer frames in buffer pool

Table 3.3: Parameters of the Resource Manager.

4. EXPERIMENTS AND RESULTS

In this section, we present performance results for the priority buffer management algorithms described earlier. In [Care89] it was shown that results for two priority levels can be generalized to multiple priority levels, so we consider just two priority levels: "low" versus "high" priority. In addition, our workload consists of three types of single-query transactions: "looping", "random reaccess (RR)", and "scanning" transactions. As their names suggest, looping transactions (such as nested-loops joins) reaccess some of their pages sequentially a number of times, RR transactions (such as hash joins) randomly reaccess some of their pages, and scanning transactions (such as clustered-index selections) touch each page just once. Looping transactions and scanning transactions represent two ends of the spectrum of buffer access characteristics typical in relational databases, while RR transactions represent the middle.

4.1. Performance Metrics

As discussed earlier, we use an open queuing system to model the DBMS. Our primary performance metric will be the average response time ratio (RTRatio) for transactions at each priority level. We define the RTRatio of a transaction as the ratio of the actual response time of the transaction to its estimated response time on an unloaded system with an infinitely large buffer pool. A transaction's response time is computed by subtracting the time at which the transaction commits from the time at which it was submitted to the DBMS. The response time of the transaction in an unloaded system is estimated by summing the CPU requirements associated with the page accesses of the transaction and by assuming one I/O per distinct page referenced by the transaction.⁹ That is, only one I/O is assumed for a page, whether it is touched just once by a transaction or accessed repeatedly. The RTRatio of a transaction, then, reflects the effects of the finite size of the buffer pool and the presence of competing transactions on the response time of the transaction. As the load on the system increases, contention for buffers causes increased I/O (and an increase in the time spent waiting outside the system) while contention for disks and CPUs causes increased disk and CPU waiting times. These factors tend to increase the RTRatio of transactions. On the other hand, if there is significant data sharing, the RTRatio of a transaction would tend to be reduced because part of the transaction's read and write sets would already be in main

⁹The cost of writing dirty data to disk is not included in this sum, since the asynchronous write engine operates independently of transactions.

memory. When the workload consists of a mix of transactions with different data access patterns and different sizes, the RTRatio provides a performance measure that is equally valid for all transactions, independent of the transaction mix.

In all the experiments described here, low-priority transactions are assumed to be running in the background, and we investigate the performance impact of the arrivals of foreground high-priority transactions on the system. The arrival rate of high-priority transactions is thus varied while keeping the arrival rate of low-priority transactions fixed. As shown in [Care89], a priority-oriented DBMS can remain stable for high-priority transactions long after the combined arrival rate has become high enough to make the system unstable for low-priority transactions. Consequently, we present response time results for low-priority and high-priority transactions separately for each experiment. Each simulation was run long enough so that the 95% confidence intervals of the RTRatios of the high-priority transactions were within 10% of the mean.

4.2. Base Parameter Settings

We first present the parameters that were used in our base experiment. As subsequent experiments are discussed, we describe the variations in the parameters for each experiment. The workload-independent parameters and the mix of transactions for each priority level in the base experiment are listed in Table 4.1. Details of the parameters for each type of transaction are presented in Table 4.2.

Workload-Independent Parameters: The database is modeled as a collection of 50 relations. We use five different relation sizes in our experiments — 1000 pages, 500 pages, 5 pages, 4 pages, and 3 pages — with the database containing 10 relations of each size. The 1000-page and 500-page relations each have a clustered index available, while the smaller relations are not indexed. There are four CPUs and four disks in the system. Each disk has 1000 tracks, and the sum of the rotational latency and the transfer time per disk access is 15 milliseconds. The

<i>NumRelations</i>	50
<i>RelSize_i</i>	1000, 500, 5, 4, 3 pages (10 relations of each size)
<i>Indexed_i</i>	YES (1000- & 500-page relations) NO (3-, 4-, and 5-page relations)
<i>IndexType_i</i>	Clustered (1000- & 500-page relations)
<i>Fanout_i</i>	20 (1000- & 500-page relations)
<i>NumCPUs</i>	4
<i>NumDisks</i>	4
<i>NumTracks</i>	1000
<i>DiskConst</i>	15 ms
<i>SeekFactor</i>	0.6 ms
<i>NumBuffers</i>	50
<i>Transaction Mix</i>	50% looping, 50% scanning
<i>Arrival Rate</i>	0-15 trans/sec (high priority) 5 trans/sec (low priority)

Table 4.1: Base Parameter Settings.

prioritized elevator disk scheduler has two priority queues, one for each priority level. The factor relating seek distance to seek time is 0.6 milliseconds, so the expected disk access time is between 15 and 30 milliseconds.

As stated in Section 1, the operating region of greatest interest to us is when the combined buffer requirements of all transactions exceeds the capacity of the buffer pool. In order to simulate the behavior of the system in this region of operation without incurring excessive simulation costs, we kept the buffer pool relatively small in our experiments. Thus, there are 50 buffer frames in the buffer pool in the base experiment. One point that should be noted here is that, from a performance perspective, it is not the actual size of the buffer pool that is most significant. Instead, two ratios are more important: the ratio of the combined buffering requirements of concurrent transactions to the size of the buffer pool¹⁰, and the ratio of the size of the buffer pool to the size of the database. For this study, we vary the first of these ratios by varying the arrival rate of high-priority transactions in all our experiments. We study the effects of varying the second ratio in Experiment 2 by changing the database size.

Workload Parameter Settings: The workload for the base experiment consists of two types of transactions. Looping transactions consist of select-joins, with the result of a selection using a clustered index on a 500-page outer relation being joined to a smaller inner relation. The selectivity of the outer relation selection varies uniformly between 0.5% and 1.5%. The inner relation is chosen uniformly from among the 30 small relations of sizes between 3 and 5 pages. *Page Accesses* is the expected number of page accesses for the transaction, with repeated references being counted as one access each time.¹¹ Scanning transactions are

Parameter	Looping	Scanning
<i>TransType</i>	Select-join	Scan
<i>JoinMethod</i>	Nested Loops	-
<i>RelSize₁ (outer)</i>	500-page	1000-page
<i>RelSize₂ (inner)</i>	3-5 pages	-
<i>AccessPath₁</i>	Cl. Idx. Scan	Cl. Idx. Scan
<i>AccessPath₂</i>	Seq. Scan	-
<i>Selectivity₁</i>	1%	1%
<i>IndexPageCPU</i>	4 ms	4 ms
<i>DataPageCPU</i>	4 ms	4 ms
<i>Page Accesses</i>	43	13
<i>Locality Set Sizes (index, outer, inner)</i>	1, 1, 3-5	1, 1
<i>Repl. Policies</i>	MRU, MRU, MRU	MRU, MRU, -
<i>Fixing Requirements</i>	3	2

Table 4.2: Workload Parameter Settings.

¹⁰Here the term "buffering requirements" of a transaction refers to the optimum number of buffers needed by the transaction.

¹¹For example, for a 1% select of a 500-page outer relation (using an index with three levels) followed by a nested-loops join with a 4-page inner relation, 43 pages are accessed: 3 to traverse the index, and then 20 pairs of (outer, inner) pages. In our model of a nested-loops join, we unfix the outer page when we unfix each inner page in order to have fre-

clustered-index scans on a 1000-page relation, with the selectivity varying uniformly between 0.5% and 1.5%.¹² For each relation accessed by a transaction, the actual relation accessed is chosen uniformly from among the 10 relations of that size. The locality set sizes and replacement policies for Priority-DBMIN are also listed in Table 4.2, as are the fixing requirements for the two types of transactions. Note that the pages of the inner relations in looping transactions will be "favored" in the Priority-Hints algorithm. The parameters for the base experiment were chosen to provide a moderate background load. (In the absence of any high-priority transactions, and with a low-priority arrival rate of 5 transactions/second, the disk utilization was in the range of 40%-50% for the different algorithms and the CPU utilization was below 20%.)

4.3. Buffer Management Without Priority

We precede the description of the base experiment by the analysis of an experiment in which priority does not affect buffer management decisions. This will help us to separate the effects of buffer management algorithms *per se* on system performance from the impact of using priority in subsequent experiments. The workload in this experiment is the same as that described in Tables 4.1 and 4.2 (except that foreground and background transactions have the same priority). In Figure 4.1, we show the RTRatios of the foreground transactions when Priority-LRU, Priority-DBMIN, and Priority-Hints are used.¹³ In order to understand the behavior of Priority-Hints relative to Priority-LRU, we also present the results for an algorithm we call "Global-Hints." Global-Hints differs from Priority-Hints in that when a buffer miss occurs and there are no free pages, the buffer manager searches the buffer pool for the globally most-recently-unfixed page and chooses it as the replacement victim. In contrast, under the same conditions in Priority-Hints, the buffer manager attempts to find unfixed pages of lower priority to choose as replacement victims, as explained in Section 2. Finding no pages of lower priority (since all transactions have the same priority in this experiment), Priority-Hints chooses the most-recently-unfixed page of the requesting transaction as the replacement victim. Thus, Priority-Hints' replacement policy is local, while Global-Hints' replacement policy is global.

In Figure 4.1, we see that Priority-DBMIN provides the best performance. The performance of Priority-Hints is close to that of Priority-DBMIN over a wide range of arrival rates, although with Priority-Hints, the system saturates at a lower arrival rate

quent suspension-safe points (see Section 2.1). This is why there are 20 (outer,inner) page pairs, with each of five outer pages looping over the set of four inner pages.

¹²In our study of the impact of updates on priority-based buffer management, non-key attributes of the tuples selected by the clustered index scans are updated. In all other experiments, the workload consists of read-only transactions, since we are interested here primarily in buffering issues rather than concurrency control.

¹³The RTRatios of the background transactions are (as one would expect) almost exactly the same as the RTRatios of the foreground transactions, and are thus not presented here.

than with Priority-DBMIN. Finally, Global-Hints provides better performance than Priority-LRU, but is significantly worse than Priority-Hints.

Priority-DBMIN provides better performance than Priority-Hints because the admission control policy of Priority-DBMIN uses its knowledge of the optimum number of buffers required for each transaction, while Priority-Hints' admission control policy cannot distinguish between the buffer requirements of looping transactions and those of scanning transactions. Consequently, Priority-Hints allows looping transactions to enter the system even when their loops cannot be guaranteed to fit in the buffer pool. This results in a higher buffer miss ratio for Priority-Hints than for Priority-DBMIN, and causes the system to become unstable at a lower arrival rate.

When we move to the Global-Hints algorithm from Priority-Hints, the local search for MRU replacement victims is replaced by a global search. This change results in a significant performance degradation for the following reason: In Priority-Hints, for looping transactions, once a transaction is able to obtain enough buffers to keep its loop (the inner relation in the nested-loops join) in memory, it proceeds quickly since it never has to give up any buffers. (Recall that no transaction can steal buffers from any other transaction in Priority-Hints for this workload.) In contrast, Global-Hints steals buffers indiscriminately from all transactions, often depriving looping transactions that have their entire working set in memory of some of their favored buffers. This causes a significant increase in disk activity for Global-Hints as compared to Priority-Hints. As a result, the system becomes unstable for Global-Hints at a foreground arrival rate of approximately 5 transactions/second, while Priority-Hints keeps the system stable until an arrival rate of about 7.5 transactions/second.

Finally, the difference between the curves for Global-Hints and Priority-LRU is caused by two features of Global-Hints. Firstly, MRU is a better search strategy for buffer replacement than LRU when the workload contains looping transactions. Secondly, Global-Hints frees normal pages as soon as it unfixes them; Priority-LRU does not. This results in favored pages being chosen as replacement victims more frequently in Priority-LRU than in Global-Hints.

This experiment shows that even in the absence of priority, Priority-Hints' local MRU replacement strategy for favored pages provides significantly better performance than Priority-LRU for our base workload, and matches the performance of Priority-DBMIN over a wide range of arrival rates. It also isolates the relative impact of the following buffer management features: admission control, which causes the difference between Priority-DBMIN and Priority-Hints; local vs. global search for replacement victims, which causes the difference between Priority-Hints and Global-Hints; and the use of MRU vs. LRU search strategies in a looping workload, which is the major factor causing the gap between the curves for Global-MRU and Priority-LRU. We can now begin our investigation of the performance of the system when the workload consists of transactions of different priority levels.

4.4. Experiment 1: The Base Experiment

In this experiment, we study the impact of using priority in buffer management for our base workload. Figure 4.2 shows the RTRatios for high-priority transactions for five buffer management algorithms: Priority-DBMIN, Priority-Hints, Priority-LRU, Global-Hints, and Global-LRU. RTRatios for low-priority transactions are shown in Figure 4.3. We explain the results of the base experiment in detail in order to provide insights into the important issues involved; this will allow us to present the results of subsequent experiments more briefly.

Comparing the curves for each algorithm in Figure 4.1 with the corresponding curves in Figure 4.2, we see that we have achieved the primary goal of priority scheduling, which is to provide a higher level of performance for high-priority transactions. For example, the system remains stable for a foreground arrival rate of up to 12 transactions/second in Figure 4.2 for Priority-Hints, while the system saturates at a foreground arrival rate of about 6 transactions/second in Figure 4.1 for the same algorithm. Of course, there is a corresponding price which is paid by low-priority transactions, as is made clear by comparing Figures 4.1 and 4.3. A secondary goal of priority scheduling is to minimize the penalty imposed on low-priority transactions; distinctions between the different algorithms in this respect will become clear as we describe subsequent experiments.

From Figure 4.2, we see that the behavior of Priority-Hints for high-priority transactions is very close to that of Priority-DBMIN, and both are superior to the other three algorithms. An interesting feature of the behavior of these two algorithms is the tradeoff between the time spent by transactions waiting outside the system in Priority-DBMIN and the time spent inside the system competing for resources in Priority-Hints. Priority-DBMIN's conservative admission policy causes the transactions' mean time spent waiting outside the system to increase more as the load on the system is increased than does Priority-Hints' liberal admission policy. However, since there are more transactions within the system in Priority-Hints than in Priority-DBMIN, the buffer miss ratios and the mean waiting times at the disks are higher for Priority-Hints than for Priority-DBMIN.¹⁴ This tradeoff will be referred to again in the following sections, where we will refer to it as the "conservative-liberal (C-L)" tradeoff. In Figure 4.4, we present the mean number of transactions (both total and high-priority) that are allowed to run concurrently by the two algorithms. Figure 4.5 shows the mean normalized DiskTime and the mean normalized OutWaitTime for high-priority transactions for the two algorithms. DiskTime is the time spent by a transaction at the disks, including the actual I/O service time and time spent waiting for disk service; OutWaitTime is the time spent by a transaction waiting outside the system. DiskTime and OutWaitTime for a transaction are normalized by dividing each of them by the expected response time of the transaction in an unloaded system.

¹⁴Contention for the CPU is not a significant factor here.

Figure 4.4 indicates that Priority-Hints allows up to 18 high-priority transactions into the system, while Priority-DBMIN limits the number of concurrent high-priority transactions to 12. Figure 4.5 shows the consequences of this. The DiskTime curves reflect the relative disk contention in the two algorithms, and Priority-DBMIN is the clear winner in limiting disk contention; both the average buffer miss ratio and the average disk waiting times per transaction are higher in Priority-Hints than in Priority-DBMIN. In contrast, Priority-Hints is the winner in limiting OutWaitTime up to an arrival rate of up to 11 transactions/second. As the load is increased beyond this, however, the disk utilization nears 100% for Priority-Hints, causing the system to saturate and Priority-Hints' OutWaitTime to exceed that of Priority-DBMIN. Priority-DBMIN's conservative admission control policy enables it to keep the system stable for arrival rates of up to 13 transactions/second.

Figure 4.4 also reveals an important point about the range of operation of greatest interest to us. There is a gap between the total number of concurrent transactions and the number of high-priority transactions for most arrival rates shown; this gap corresponds to the number of low-priority transactions running in the system. When the curves for the total number of transactions in the system flatten out, the buffer pool has become fully utilized, but Figure 4.4 shows that there are still significant numbers of low-priority transactions running in the system. It should be clear that priority-based buffer replacement policies will be most useful in this range of operation, since buffers owned by low-priority transactions can be stolen by high-priority transactions. In Figure 4.4 we also see that as the arrival rate of high-priority transactions increases, low-priority transactions are gradually displaced by high-priority transactions (due to the use of priority-based admission control policies) until finally only high-priority transactions remain active. Thus, both priority-based admission control and priority-based buffer replacement have an important role in determining performance over a fairly wide range of arrival rates.

Figure 4.2 shows us that at low loads, all the algorithms provide similar levels of performance to high-priority transactions. As the high-priority load increases, the curves for Global-LRU, Global-Hints, and Priority-LRU soon branch away from the curves for Priority-Hints and Priority-DBMIN. Global-LRU performs worst of all: it does not distinguish between transactions of high and low priorities, and it uses the LRU criterion for replacement. As for Global-Hints, its ability to distinguish between favored and normal pages (and its use of MRU) actually allows it to perform better than Priority-LRU at very low high-priority loads. As the load is increased, however, Priority-LRU's protection of high-priority buffers begins to have a greater impact, since more of the buffers are now owned by high-priority transactions. The RTRatio of high-priority transactions remains lower for Priority-LRU than for Global-Hints until the system becomes unstable for both algorithms at a high-priority arrival rate of approximately 8 transactions/second.

In Figure 4.3, the curves for Priority-Hints and Priority-DBMIN are fairly close. This is because the criteria used for suspending low-priority transactions differ in these two algorithms, making the C-L tradeoff more even for low-priority

transactions than it is for high-priority transactions. (High-priority transactions cannot be suspended in either algorithm.) Priority-DBMIN suspends a low-priority transaction immediately when one of its unfixed buffers is required by a high-priority transaction. In contrast, Priority-Hints does not suspend low-priority transactions as frequently as Priority-DBMIN does. It allows them to continue execution as long as their "fixing requirements" can be satisfied; of course, this increases the low-priority load on the disks. As long as there is sufficient disk capacity to handle this increased low-priority load in Priority-Hints, the two algorithms provide similar performance for low-priority transactions.

Figure 4.3 also shows that there is very little difference in the performance provided by the LRU algorithms (Global-LRU and Priority-LRU) for low-priority transactions. Priority-LRU steals buffers from low-priority transactions in preference to depriving high-priority transactions of their buffers, so one might expect Priority-LRU to provide worse performance for low-priority transactions. Recall, however, that the CPUs and the disks use priority scheduling in these experiments; also, in the range of arrival rates for which the system is stable for low-priority transactions, there are relatively few high-priority transactions in the system. The globally least-recently-used buffer is therefore quite likely to belong to a low-priority transaction rather than to a high-priority transaction. This is why the curves for Priority-LRU and Global-LRU are so close to each other in Figure 4.3.

Low-priority transactions perform better under the Global-Hints algorithm than under the two LRU algorithms in Figure 4.3 because their buffer miss ratios are lower due to the use of MRU. The reason that Global-Hints performs worse than Priority-Hints there, even for low-priority transactions, is again related to the use of priority at the CPUs and at the disks. As the arrival rate of high-priority transactions increases, Global-Hints hurts high-priority transactions more than Priority-Hints does (since Global-Hints ignores priority in choosing replacement victims). Consequently, more and more of the system's disk capacity is used to satisfy high-priority transactions in Global-Hints. This makes the disk waiting times of low-priority transactions higher, causing the system to become unstable for low-priority transactions at a lower load in Global-Hints than in Priority-Hints.

The base experiment confirms the result of [Care89] for our mixed workload: the use of priority in buffer management is a clear win (independent of the algorithm) if the response time of high-priority transactions is the main criterion of system performance. However, the conclusions are more mixed for low-priority transactions: their performance may be worse for some priority-based buffer management algorithms (e.g., Priority-LRU) than for algorithms that do not consider priority in replacement decisions (Global-Hints). The base experiment also shows that for both low-priority and high-priority transactions, the performance provided by Priority-Hints is significantly better than the performance provided by Priority-LRU, and that Priority-Hints performs almost as well as Priority-DBMIN for both priority levels. In subsequent experiments, we limit ourselves to showing the relative behavior of the three priority-based algorithms.

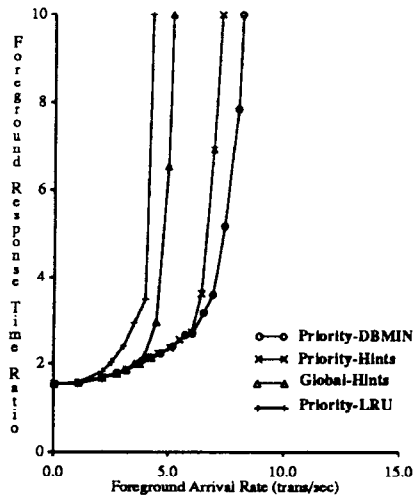


Fig. 4.1: No Priority.

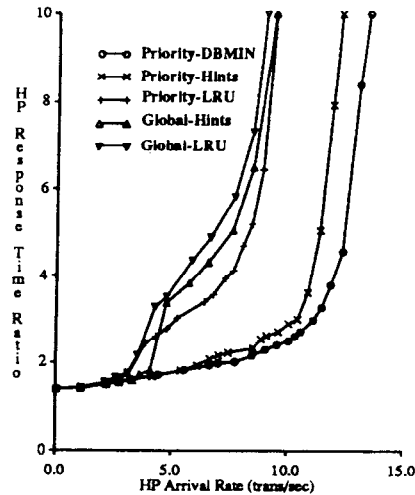


Fig. 4.2: High Priority.
(Base Experiment)

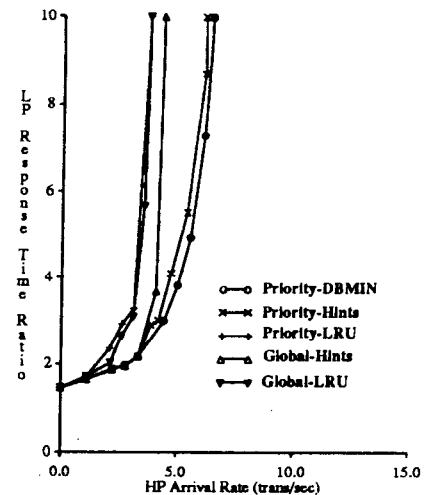


Fig. 4.3: Low Priority.
(Base Experiment)

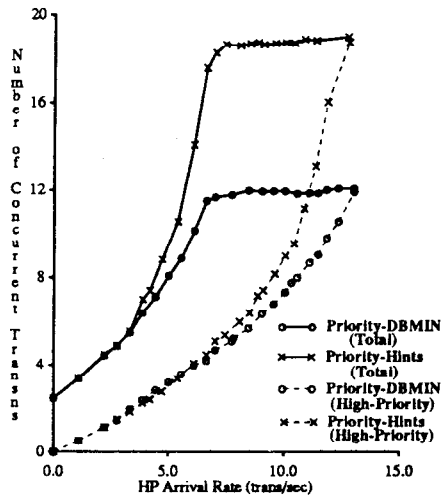


Fig. 4.4: Number of Concurrent Transactions.
(Base Experiment)

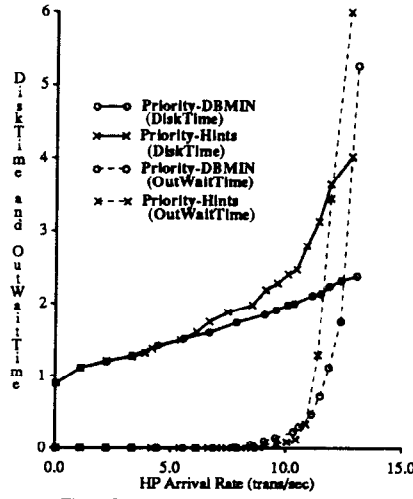


Fig. 4.5: DiskTime and OutWaitTime.
(Base Experiment)

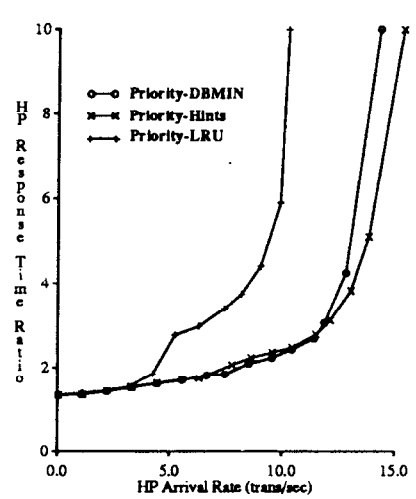


Fig. 4.6: High Priority.
(Medium Data Sharing)

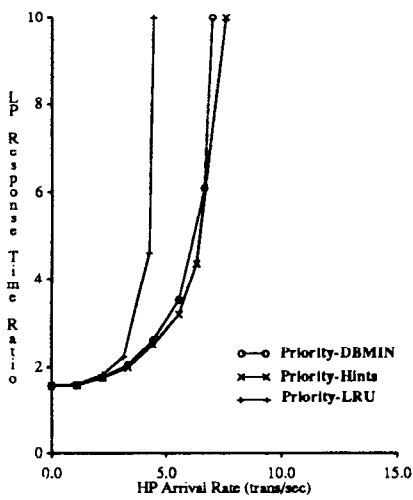


Fig. 4.7: Low Priority.
(Medium Data Sharing)

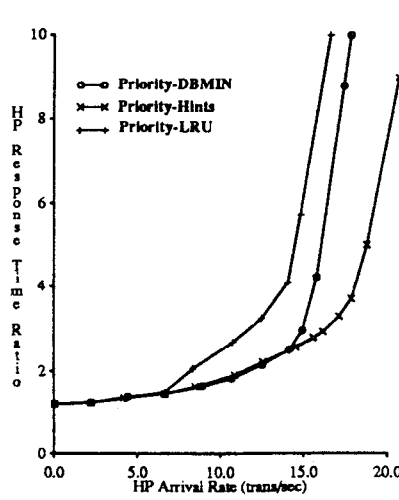


Fig. 4.8: High Priority.
(High Data Sharing)

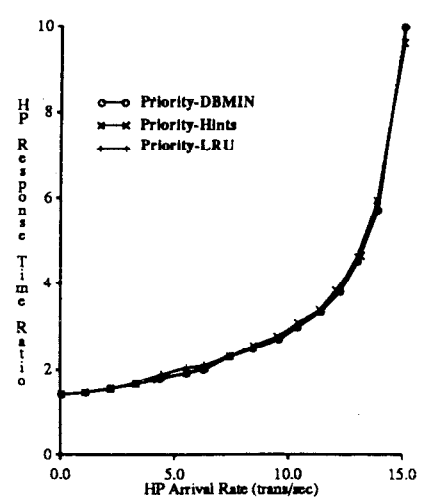


Fig. 4.9: High Priority.
(HP Load = 100% scanning)

4.5. Experiment 2: Varying Relative Buffer Pool Size

In this experiment, we reduce the size of the database while keeping all other parameters the same as in the base experiment. Thus, the ratio of the size of the buffer pool to the size of the database is higher in this experiment than in the base experiment. Increasing the relative size of the buffer pool in this manner results in increased data sharing: i.e., data brought into the buffer pool at the request of one transaction is more likely to be found in memory when it is accessed by other transactions. When data accesses are distributed uniformly over the database, the extent of data sharing is inversely proportional to the database size if all other parameters are kept fixed. In the base experiment, the database consisted of 50 relations and a total of 15,120 data pages, while the buffer pool had 50 buffers (see Table 4.1). This represents a fairly low level of data sharing (and we did not even take index pages into account when computing the sum of 15,120 pages!). We consider two levels of data-sharing in this experiment: one where there are 25 relations in the database (five relations each of 1000 pages, 500 pages, 5 pages, 4 pages, and 3 pages), and one where there are just 5 relations (one of each size). When there are 25 relations in the database, the level of data sharing is double the level of data sharing in the base experiment; when there are just 5 relations, the level of data sharing is ten times that in the base experiment.

In Figure 4.6, we present the RTRatios of high-priority transactions for the three algorithms with 25 relations in the database. Figure 4.7 shows the RTRatios of low-priority transactions for the same database. As one would expect, the performance of all three algorithms improves over their performance in the base experiment. The key difference between the trends shown in Figures 4.4 and 4.2 is the fact that Priority-Hints now provides better performance than Priority-DBMIN at high loads. As the level of data sharing increases, Priority-DBMIN's conservative admission control policy proves to be more and more harmful: it simply underutilizes resources by failing to account for the possibility of data-sharing. The same conclusion holds true for low-priority transactions, where Priority-Hints consistently performs as well as or better than Priority-DBMIN.

Figure 4.8 presents the RTRatios for high-priority transactions in the case where there are just 5 relations in the database. Priority-Hints and Priority-DBMIN provide the same level of performance for high-priority transactions until Priority-DBMIN's admission control policy causes it to block transactions unnecessarily outside the system; beyond this point, Priority-Hints is better. As for Priority-LRU, one might initially expect that all three inner relations of the nested-loops joins (a total of 12 pages) would always remain in memory, and that it should therefore perform as well as Priority-Hints. However, recall that there are now 1500 data pages (plus index pages) that are not part of inner relations. A large fraction of the page requests made to the buffer manager are for one of these 1500 pages. In Priority-Hints and Priority-DBMIN, the looping pages will remain in memory as long as there are some high-priority transactions that need them. Also, these two algorithms free normal pages as soon as they are unfixd, so it is quite likely that a free page will be available even when the load is high. In Priority-LRU, where transactions can steal buffers from other transactions of the same

priority, and where looping pages are treated just like other pages, looping pages are frequently chosen as replacement victims. This is why the curve for Priority-LRU diverges from the other two in Figure 4.8.

This experiment shows that as the relative size of the buffer pool increases, the performance of Priority-Hints improves to a greater extent than that of Priority-DBMIN. This is a consequence of the latter's conservative admission control policy. Also, both Priority-Hints and Priority-DBMIN provide better performance than Priority-LRU, even under very high data sharing.

4.6. Experiment 3: Changing the Transaction Mix

In the base experiment, the high-priority workload consisted of a mix of an equal number of looping and scanning transactions. In this experiment, we vary the proportion of looping transactions in the high-priority workload while keeping all other parameters as in the base experiment. (The low-priority workload still consists of a 50% looping, 50% scanning mix.) Figure 4.9 shows the RTRatios for high-priority transactions when the high-priority workload consists entirely of scanning transactions. RTRatios for the low-priority transactions are shown in Figure 4.10. In Figure 4.9, the curves for the three algorithms coincide almost exactly. This is not unexpected, since the high-priority workload is insensitive to buffer replacement and the admission control criteria for the three algorithms coincide for scanning transactions.¹⁵ The interesting feature of this experiment is the relative performance of low-priority transactions, as shown in Figure 4.10: here, Priority-Hints performs far better than Priority-LRU. The reason for this is that, unlike Priority-Hints, Priority-LRU does not free scanning pages as soon as they are unfixd. High-priority transactions thus keep stealing looping pages from low-priority transactions in Priority-LRU, while the more appropriate action is to choose high-priority scanning pages as replacement victims.

Figures 4.11 and 4.12 show the RTRatios for high-priority transactions and low-priority transactions, respectively, when the high-priority workload consists entirely of looping transactions. There is now a relatively larger difference between the performance of high-priority transactions for Priority-Hints and Priority-DBMIN than in the base experiment. The C-L tradeoff for high-priority transactions favors Priority-DBMIN in this experiment, since a larger proportion of the high-priority workload now benefits from conservative admission control. However, note that the tradeoff is still quite even for transactions of low priority. This is because Priority-DBMIN now prevents more scanning low-priority transactions from entering the system than Priority-Hints does, and it also suspends them more often.

¹⁵The number of buffers required to meet the fixing requirements of a scanning transaction is equal to the sum of the sizes of its locality sets; see Table 4.2.

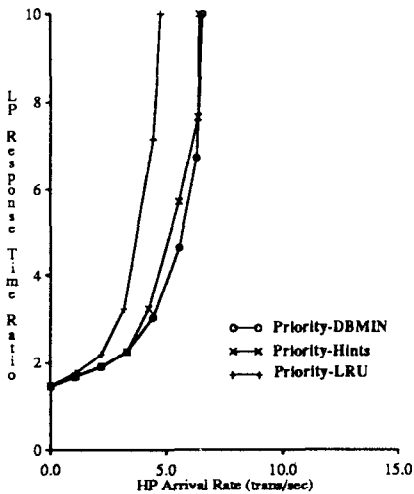


Fig. 4.10: Low Priority.
(HP Load = 100% scanning)

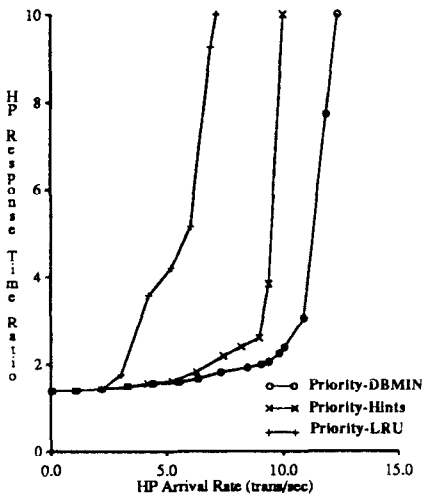


Fig. 4.11: High Priority.
(HP Load = 100% looping)

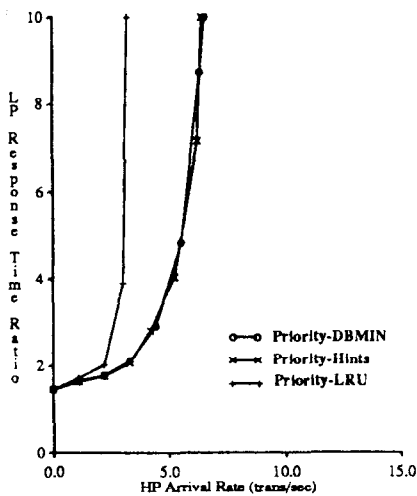


Fig. 4.12: Low Priority.
(HP Load = 100% looping)

From Experiment 3, we learn that Priority-Hints tends to perform as well as or better than Priority-LRU, independent of the proportion of looping and scanning transactions. In particular, when the low-priority load is insensitive to buffer management, Priority-Hints provides much better performance than Priority-LRU for low-priority transactions. As the proportion of looping transactions in the high-priority workload increases, Priority-DBMIN's admission control policy makes it perform better than Priority-Hints for high-priority transactions, but the two algorithms provide very similar support for low-priority transactions.

4.7. Other Experiments

We conducted a number of experiments in addition to the ones described above, in which we explored other issues related to priority-based buffer management [Jauh90]. Space considerations force us to summarize these results here.

- Random Reaccess Transactions:** We examined the performance of Priority-Hints when the workload consisted of classic hash-join queries [Shap90]. This workload, where a transaction reaccesses a number of its pages randomly, was chosen because its pattern of data accesses is similar to that of other types of hash-based joins and of non-clustered index selections. In our model of the classic hash-join, the inner relation is read into the buffer pool and a hash-table is built on its tuple identifiers in the private workspace of the transaction. Then, for each selected tuple of the outer relation, this hash table is probed for matching tuples in the inner relation. Note that this variant of the classic hash join allows pages of the inner relation to be replaced in the buffer pool, though such replacements should clearly be avoided as much as possible for performance reasons. In our experiment, each transaction consisted of a select-join, with the result of a 0.2% clustered index selection on a 1000-page relation being joined with a 5-page inner relation. Inner relation pages were "favored" in the Priority-Hints algorithm. Our experiment showed that even when the use of MRU is irrelevant, the classification of pages into "favored" and "normal" sets by Priority-Hints enables it to provide significantly better performance than Priority-LRU. Priority-DBMIN provides better performance than Priority-Hints for high-priority transactions, since it keeps the entire inner relation fixed in memory for the duration of the join. For low-priority transactions, the performance of Priority-Hints and Priority-DBMIN is very similar, for reasons closely related to those discussed in explaining Figure 4.12.

- Low-Priority Updates:** Clearly, updates result in an increased load on the disks, as dirty pages have to be written back to disk before they can be replaced in the buffer pool. As discussed in Section 1, the key issue from a priority standpoint is the effect (if any) of low-priority updates on the performance of high-priority transactions. We examined this issue by changing the low-priority workload to include scanning update transactions; a clustered-index scan was used to select tuples, and non-indexed attributes of the selected tuples were updated. The high-priority workload consisted of looping, read-only transactions. We ensured that the write sets of concurrent transactions did not intersect in this experiment, so concurrency control was not an

issue. Our experiments showed that when Priority-LRU was used, the RTRatios of the high-priority transactions increased significantly as the update probability increased, while Priority-Hints and Priority-DBMIN managed to keep their high-priority transactions almost immune to the presence of low-priority updates at the same load. Note that in this experiment, the updated pages were normal. Recall also that all three algorithms ignore dirty pages in their search for replacement victims. Priority-Hints and Priority-DBMIN place dirty normal pages in the dirty list as soon as they are unfixd, just as they place clean normal pages in the free list at unfix time. In contrast, Priority-LRU does not free any pages until transaction commit time unless all possible replacement victims are dirty, in which case a dirty page is synchronously written out to disk to prevent deadlock as explained in Section 2.5. This causes dirty normal pages to accumulate in the buffer pool in Priority-LRU.

- **Changing System CPU and I/O Capacity:** We varied the number of CPUs and disks while keeping the buffer pool size fixed in order to understand the impact of different resource capacities on the relative performance provided by the three buffer management algorithms. We found that the performance differences between Priority-DBMIN and Priority-Hints occur as a consequence of higher disk activity when the latter policy is used; if the workload is CPU-bound, these two algorithms will provide similar performance. The results also confirmed that Priority-Hints' performance is close to that of Priority-DBMIN, and superior to that of Priority-LRU, independent of the CPU, I/O and main memory capacities of the system.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we have shown that buffer management can have a very significant impact on the performance of a priority-oriented database system, especially when the unpredictability of the workload forces the system to operate in regions where the total buffer requirements of the concurrent transactions exceed the system's buffer capacity. We have introduced a new buffer management algorithm, called Priority-Hints, that uses page-level information provided by the database access methods to make priority-based buffer management decisions. The performance of our new algorithm has been compared to the performance of Priority-LRU and Priority-DBMIN, two algorithms proposed earlier for priority-based buffer management. A number of performance insights have been obtained as a result of simulation experiments. Priority-Hints was shown to perform better than Priority-LRU for all of the workloads considered here. For most workloads, Priority-Hints performed almost as well as Priority-DBMIN; for some workloads with data sharing, Priority-Hints actually provided better performance than Priority-DBMIN. Even when the workload consisted of transactions of equal priority, Priority-Hints performed significantly better than Priority-LRU and almost as well as Priority-DBMIN.

These results are significant for several reasons: First, in previous studies [Chou85, Care89] it has been shown that DBMIN-like approaches to buffer management provide better performance than approaches based on simple strategies such as LRU. Still, most existing database systems continue to use LRU-based

approaches because they do not require as much information as DBMIN does. Second, the type of information used by the buffer manager in Priority-Hints is already being provided to buffer managers in existing database systems [Teng84, Haas90]. Our algorithm has the advantages of both DBMIN-based approaches and LRU-based approaches; it provides good performance while requiring little information. Finally, we have also shown that Priority-Hints adapts itself dynamically as data sharing increases, while Priority-DBMIN's more static approach to buffer allocation can cause its performance to suffer in the presence of data sharing.

We plan to continue our work in the area of priority-based DBMS scheduling. For example, we will extend our algorithms to workloads consisting of multi-query transactions. We also intend to study the problem of concurrency control conflicts in a priority-oriented DBMS. Finally, we plan to extend our performance study to the real-time context, where the workloads contain transactions with deadlines and importance levels.

REFERENCES

- [Abbo88] Abbott, R., and Garcia-Molina, H., "Scheduling Real-Time Transactions: A Performance Evaluation," *Proc. 14th VLDB Conf.*, Los Angeles, CA, Aug. 1988.
- [Abbo89] Abbott, R., and Garcia-Molina, H., "Scheduling Real-Time Transactions with Disk Resident Data," *Proc. 15th VLDB Conf.*, Amsterdam, Aug. 1989.
- [Bitt88] Bitton, D., and Gray, J., "Disk Shadowing," *Proc. 14th VLDB Conf.*, Los Angeles, CA, Aug. 1988.
- [Care89] Carey, M. J., Jauhari, R., and Livny, M., "Priority in DBMS Resource Scheduling," *Proc. 15th VLDB Conf.*, Amsterdam, Aug. 1989.
- [Chou85] Chou, H-T., and DeWitt, D., "An Evaluation of Buffer Management Strategies for Relational Database Systems," *Proc. 11th VLDB Conf.*, Stockholm, Sweden, Aug. 1985.
- [Effe84] Effelsberg, W., and Haerder, T., "Principles of Database Buffer Management," *ACM TODS* 9(4), Dec. 1984.
- [Haas90] Haas, L., et al, "Starburst Mid-Flight: As the Dust Clears," *IEEE Trans. on Knowledge and Data Eng.*, March 1990.
- [Hari90] Haritsa, J., Carey, M. J., and Livny, M., "On Being Optimistic About Real-Time Constraints," *ACM PODS Conf.*, Nashville, TN, April 1990.
- [Jauh90] Jauhari, R., Carey, M. J., and Livny, M., *Priority-Hints: An Algorithm for Priority-Based Buffer Management*, Tech. Rep. # 911, Comp. Sci. Dept., Univ. of Wisconsin-Madison, Feb. 1990.
- [Livn89] Livny, M., *DeNet User's Guide*, Version 1.5, Computer Sciences Dept., Univ. of Wisconsin, Madison, 1989.
- [SIGM88] *SIGMOD Record* 17(1), Special Issue on Real-Time Data Base Systems, S. Son, ed., March 1988.
- [Sacc86] Sacco, G.M., and Schkolnick, M., "Buffer Management in Relational Database Systems," *ACM TODS*, 11(4), Dec. 1986.
- [Shap86] Shapiro, L.D., "Join Processing in Database Systems With Large Main Memories," *ACM TODS*, 11(3), Sept. 1986.
- [Teng84] Teng, J., and Gumaer, R. A., "Managing IBM Database 2 buffers to maximize performance," *IBM Sys. J.* 23(2), 1984.