

Exploiting Inheritance and Structure Semantics for Effective Clustering and Buffering in an Object-Oriented DBMS

Ellis E. Chang, Randy H. Katz

Computer Science Division
Electrical Engineering and Computer Science Department
University of California, Berkeley
Berkeley, CA 94720

ABSTRACT: Object-oriented databases provide new kinds of data semantics in terms of inheritance and structural relationships. This paper examines how to use these additional semantics to obtain more effective object buffering and clustering. We use the information collected from real-world object-oriented applications, the Berkeley CAD Group's OCT design tools, as the basis for a simulation model with which to investigate alternative buffering and clustering strategies. Observing from our measurements that real CAD applications exhibit high data read to write ratios, we propose a run-time clustering algorithm whose initial evaluation indicates that system response time can be improved by a factor of 200% when the read/write ratio is high. We have also found it useful to limit the amount of I/O allowed to the clustering algorithm as it examines candidate pages for clustering at run-time. Basically, there is little performance distinction between limiting reclustering to a few I/Os or many, so a low limit on I/O appears to be acceptable. We also examine, under a variety of workload assumptions, context-sensitive buffer replacement policies with alternative prefetching policies.

1. Introduction

Object-oriented concepts, developed originally in Smalltalk, have been accepted by database researchers as the basis for dealing with semantics-rich applications like CAD/CAM. The key aspect of an object-oriented system is the mapping from an application's logic into a set of abstract data types, with associated operations and attributes. In addition, information (operation/attribute definitions) can be propagated along the lattice formed by instances, types and super-types through inheritance.

These object-oriented concepts have been further enhanced to support complex object, version and correspondence management, and new ways to propagate information between instances [KATZ87, CHAN87a]. For example, in our Version Data Model, we explicitly support three structural relationships, i.e. configuration, version history and correspondence (see Figure 1.1). Note that objects are named by the triple *name[i].type*, where *name* is the object name, *i* is the version number, and *type* is its representation type.

We have proposed a new *instance-to-instance* inheritance in which information can be propagated among instances along structural relationships. For instance, there are some properties and behaviors, as well as structural relationships and constraints, that an

offspring version might wish to inherit directly from its parent version rather than its type. Consider the following example. If ALU[2].layout corresponds to ALU[3].netlist, then a new descendant of ALU[2].layout should inherit this correspondence relationship by default. Just as the relational query optimizer uses the table cardinality and indexing information to produce efficient access plans, object-oriented systems could also use knowledge of structural relationships and inheritance to improve system performance [CHAN87a].

In this paper, we are particularly interested in how clustering and buffering algorithms can use structural relationships and inheritance semantics to improve DBMS response time. For instance, if a design browser walks through multiple representations of the same design objects, clustering across correspondence is advantageous. Alternatively, if a simulation tool traverses the net list representation hierarchy, clustering along the configuration hierarchy is best. If information is frequently inherited along the version history, the system may place the object near its ancestor object to save disk space. Moreover, the buffer manager may accept *hints* from the clustering algorithm to keep candidate pages for clustering in the buffer pool, to avoid I/Os during clustering and to improve response time.

A number of object-oriented database management systems have implemented object-based clustering mechanisms [MAIE86, ATWO85, ZDON84, KIM87]. Two characteristics are common to each of these: (1) the segment is the clustering unit, and (2) user's hints are used, but only at object creation time. For instance, a user may provide a hint such as "place near object XX". The system would then try to store the target object with object XX in the same page or adjacent page. Otherwise, the first buffer page with available space is chosen even though other buffer pages may contain the target object's structural related objects, e.g. ancestor or composite objects. Since these systems do not model structural relationships as first class objects, the storage component has no structural information to exploit during the clustering process. Users' hints are the only useful semantics used by the storage component. Furthermore, physical placement hints based on instance-to-instance inheritance are not exploited by any system. Neither do these systems consider reclustering when object structures are changed.

We will describe a run-time reclustering algorithm which can improve the overall system response time by 200% when the read/write ratio is high. We have found it useful to limit the amount of I/O allowed to the clustering algorithm as it examines candidate pages for reclustering at run-time. We will also discuss how prefetching with an alternative scope of candidate pages affects response time when different buffer replacement policies are used.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.
© 1989 ACM 0-89791-317-5/89/0005/0348 \$1.50

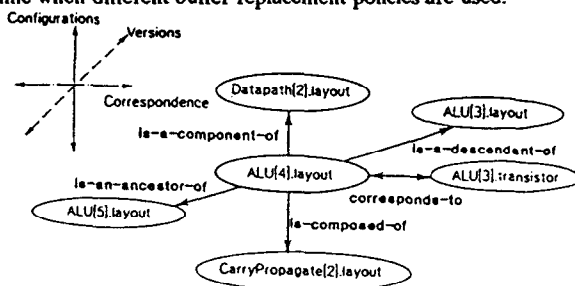


Figure 1.1 - Version Data Model

We have instrumented an object-oriented Data Manager OCT, developed by the UC-Berkeley CAD group [HARR86], and have collected the access pattern information of more than ten CAD tools running on top of OCT. About 5000 tool invocations representing approximately 400 hours of design work are recorded. We have observed very high data read to write ratios from the OCT tools environment in the measurement results, which implies that dynamic clustering and context-sensitive buffering can be very useful in object-oriented applications.

The rest of the paper is organized as follows: Section 2 discusses how structural relationships and inheritance can be used in various clustering and buffering algorithms. Section 3 reports the OCT tools' access pattern, read/write ratio, I/O rate and structure densities which are subsequently used to define the workload in the simulation model. Section 4 describes the simulation model, constructed in a modeling language called Performance Analyst's Workbench System (PAWS) [PAWS83] to evaluate the performance impact of different algorithms. The simulation results are discussed in Section 5 through "Two-level factorial" analysis technique. Finally, our conclusions and future directions are given in Section 6.

2. Effective Clustering and Buffering

Engineering design, manufacturing, and CASE applications demand high performance object management systems. These applications utilize structural relationships and inheritance mechanisms to effectively model complex environments. Only if conventional relational database management systems support these modeling primitives could the storage component take advantage of them at run-time.

In addition, these complex applications frequently perform materializations of an object hierarchy. Loading a large object hierarchy into memory becomes the bottleneck. This section describes how to obtain better bandwidth by exploiting inheritance and structural semantics in buffering and clustering for object-oriented applications.

2.1. Smart Clustering

The objective of clustering is to place frequently co-referenced objects near each other in physical memory. For *static* clustering, the system is quiesced, and the database administrator decides on a partitioning of objects. *Static* clustering is not effective for applications such as manufacturing which require high availability. *Dynamic* clustering, on the other hand, is done at run-time when concurrent accesses are permitted. Although clustering on object creation or update may degrade the response time of writers, such degradation to writers can be offset by a large improvement of readers' response time as the simulation results in Section 5 show. *Dynamic* clustering, therefore, becomes very attractive in object-oriented applications where reads dominate writes (Refer to Section 3).

The inputs to the clustering algorithm can be user's hints such as "access by configuration", the interobject access frequencies, and the characteristics of inherited attributes. The user's hints are registered into the system through a procedural interface. The interobject access frequencies are inherited from the type at object creation time. For instance, in CASE applications, the run-time debugger frequently navigates from the object code to the source code and not in the other direction. Such information can be predefined at type creation time and used by all instances.

The clustering algorithm (the full algorithm is presented in [CHAN87a]) chooses an initial placement for each newly created instance based on which of the instance's relationships is most frequently traversed. This frequency information is available in the corresponding data type and is inherited by the newly created

instance. These frequencies may be affected by the implementation of the instance's inherited attributes. For inherited attributes, the clustering algorithm uses an additional set of cost formulas to choose between implementation by copy versus by reference. If the implementation is by copy, the system needs to maintain the consistency among all duplicated attribute values at run-time. Furthermore, the augmented access frequencies may change the initial placement of the instance.

The following set of parameters is used to control the clustering algorithm:

- (1) *Candidate page pool.* Since the information used by the clustering algorithm is on instance basis, the clustering algorithm needs to retrieve the physical page in order to get the corresponding information for clustering decision. Notice that the free space map in the page table is not sufficient to determine the potential candidate page. When looking for a candidate page for placement, the clustering algorithm may use only the pages available in the buffer pool, avoiding any I/Os during the clustering process. Or, the algorithm may search a limited number of pages on disk. Alternatively, if the I/O limit is infinite, then the algorithm may use the entire database as the candidate page pool. The candidate page pool can be `Within_Buffer`, `With_IO_limit`, or `Within_DB`.
- (2) *Page splitting policy.* When the preferred candidate is full, the storage manager must either split the page to make room for the new object, or choose the next best candidate page which has space for the new object. The page is split if the expected access cost resulting from the page-split is better than the cost of putting the new object in the next best candidate page. Otherwise, the next candidate page is examined, and the decision process recurses if there is insufficient room on the chosen page.

The problem of estimating the cost resulting from the page-split is similar to the Graph Partitioning Problem. However, this is known to be NP-complete and may not be suitable for a run-time clustering algorithm. [CHAN87a] identifies a greedy algorithm that partitions the nodes of the inheritance-dependency graph into two subsets that can fit into a page individually. At the same time, the greedy algorithm tries to minimize the total accessing cost. Because the algorithm does not try to find the optimal partition and only scans through the set of arcs once, the total running time is guaranteed to be linear. If the degree of connection and the number of nodes are small, the complexity of the page splitting algorithms should have no major impact on the overall system response time. Therefore, the page splitting policy can be `No_Splitting`, `Linear_Split`, or `NP_Split`.

- (3) *User Hints Policy.* Many computer-aided design applications make frequent use of configuration relationships. However, most inheritance references are along version history relationships, since a descendant version typically obtains information from its ancestor. If users can make these reference information known to the system, the system can cluster objects based on the characteristics of the application. The User hints policy, can be designated as `User_Hints` or `No_Hints` in order to discover the effectiveness of these user hints.

2.2. Smart Buffer Replacement

The most common operations of object-oriented tools are navigation along the structural relationships and simple retrieval of design objects. In general, object-oriented tools have fairly static access patterns. Unfortunately, these are ignored by most database systems.

To obtain better response time, the buffer manager must determine prefetching and buffer replacement strategies by exploiting knowledge about the structural and inheritance relationships. Response time is improved if appropriate objects can be prefetched before they are needed, or if related objects can be kept in the buffer pool even if the relationships span disk pages.

The implementation overview of an object-oriented buffer manager is presented here. (Refer to [CHAN87a] for more detailed algorithms) The least sophisticated buffer manager uses a simple LRU buffer replacement policy and attempts no prefetching. A more sophisticated approach uses a priority scheme which replaces the lowest priority pages first. The key challenge is to use the semantics of the interrelationships among objects on the buffered pages and hints about the access patterns to set the priorities intelligently. Frequently accessed pages have their priority increased. Infrequently accessed pages have their priority reduced, but this may be modified by their interrelationships with other pages, especially if those are frequently accessed. Whenever an object is accessed, its related pages (these pages containing its component objects and its inherited attributes) might be in the buffer pool already. The traditional LRU buffer replacement algorithm could easily choose these pages to be replaced, and thus cause extra I/Os to bring them in later on.

Another way to obtain good response time is to take advantage of prefetching. At the beginning of an interaction with the database, the users provide the buffer manager with access hints, such as "my primary access is via configuration relationships". This information influences the buffer manager's prefetch strategy. Touching an object causes the page containing it and the pages containing its immediate component objects to be brought into the buffer pool and given the same high priority. This achieves extremely good performance for applications that walk the configuration hierarchy. Similar prefetch hints can be used to obtain a version object, its immediate ancestor, and its immediate descendents. Also, correspondence relationships can be used to obtain all objects corresponding to the one being accessed. Inheritance is treated in a similar fashion for determining prefetch groups.

The following set of parameters is used to control the buffering algorithm:

- (1) Buffer Pool Size.
- (2) Buffer Replacement Policy. Three policies are examined: Context-sensitive, LRU and Random.
- (3) Prefetch Policy. The candidate pages for prefetching can be constrained to either the buffer pool or the entire database. Notice that prefetching within buffer pool does not create any extra logical I/Os. However, it will cause the buffer priority to be adjusted to the requesting applications.

2.3. Experiments to be evaluated

These control parameters are used in the simulation model described in Section 4. In this paper, we only report the results of the following experiments:

- (1) Alternative clustering policies with varying page splitting policies.
- (2) Alternative buffer replacement policies with varying prefetching policies.

For more detailed results, refer to [CHAN89].

3. Object-Oriented Application Access Pattern

We instrumented OCT, an object-oriented data manager, interviewed several OCT's users who are VLSI CAD tool developers, and collected access patterns from real user environments. OCT has more than 20 client applications, and it is being used heavily in the Berkeley CAD community. We felt that OCT would give us a very good starting point in understanding the access pattern of object-oriented applications even though it provides only a subset of object-oriented concepts. The following section reports on our findings.

OCT, data manager for VLSI/CAD applications [HARR86], supports a set of primitive object types which are used frequently by VLSI CAD tools and arbitrary attachments among these objects. Figure 3.1, for example, shows a net attached to a facet which is a basic design unit, four terminals attached to the net, and three paths attached to various terminals. The run-time navigation path is represented by the arc direction from one object to another. Every attachment creates a link between two objects. These links are bidirectional and provide basic composition hierarchy information. However, users are responsible for maintaining the legal attachment among objects, and the system does not provide any structure validation. Further, OCT supports no explicit inheritance mechanisms. situation where OCT uses inheritance, is not available to OCT users.

3.1. Information Collected

For each OCT tool invocation, we recorded the following information:

- (1) The tool Identifier, such as "SPARCS" or "VEM", helps us understand the run-time behavior of an individual application and also allows us to relate the tools' functionalities to their overall run-time behavior. For instance, we like to know whether all the placement and routing tools have similar read/write ratios.
- (2) Read and write activities. When objects are retrieved through "attachment" links, they are recorded as *structure read*. Any links created between objects via "attachment" are viewed as *structure write*. Remaining read or write operations are viewed as *simple read* or *simple write*. We also recorded the object type information for every read and write operation.

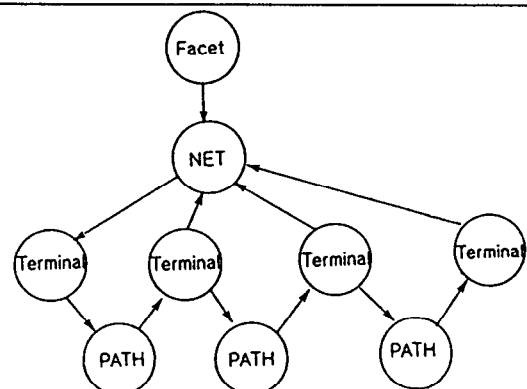


Figure 3.1 – MOSAICO Access Pattern

- (3) Session time. A session is defined as the time interval between `octBegin()` and `octEnd()`. With this, we can measure the I/O rate per session and understand the correlation between the session time and the applications' functionalities.
- (4) Fan-out of upward and downward structural access. Reading a composite object may likely cause the retrieval of a subset of its component objects. Similarly, reading a component object may cause the return of several of its composite objects. For example, a netlist simulator would only navigate along the `<cell>`, `<net>`, and `<segment>` path, leaving the remaining component objects attached to `<cell>` totally untouched. In other words, not all of the component objects are read when an application is reading a composite object at run-time.

3.2. R/W ratio

For every tool invocation, we collected the number of structure read, structure write, simple read, and simple write operations, all of which are at the logical level. We define the read/write ratio of every tool invocation to be the total number of structure reads and simple reads divided by the total number of structure writes and simple writes. Figure 3.2 shows the read/write ratio of nine OCT tools. VEM, a graphical editor not included in Figure 3.2, has the highest read/write ratio of 6000. The rest of the OCT tools' read/write ratios vary from 0.52 to 170. One interesting note is that different phases of the same application may have wide variations in the read/write ratio. For instance, the macro cell router MOSAICO is composed of `atlas`, `cds`, `cprep`, `PGcurrent` and `mosaico`. Figure 3.2 shows the read/write ratio within one run to vary from 0.52 to 170. This is quite unusual compared to traditional debit/credit applications, in which the read/write ratio is quite stable throughout an application. Due to this dynamic of applications' read/write ratio, the clustering algorithm must be adaptive to achieve adequate response time at different phases of an application.




Figure 3.2 – OCT Tools' Read-Write Ratio




Figure 3.3 – OCT Tool Structure Density Distribution

3.3. Structure Density

Figure 3.3 shows the downward access structure density of ten OCT tools. Although both upward and downward accesses are measured at run-time, we observe that most of the upward accesses have only one object returned. Therefore, we only report the fan-out of downward access here.

Most of the OCT tools' downward access except VEM are dominated by a low structure density (0 to 3 objects). VEM has the highest structure density, since it typically displays all the objects attached to the composite object.

3.4. Access Pattern

After interviewing several OCT tool developers, we observed the following:

- (1) Object-oriented applications perform more navigation than the ad-hoc query during run-time. For instance, in Figure 3.1, the macro cell router MOSAICO navigates along the configuration to find all the paths used by a certain net. Ad-hoc query like "select all nets having 6 or more terminals" is not needed in MOSAICO.
- (2) Certain access patterns can be eliminated if the underlying system supports integrity constraint. For example, the cell compactor SPARCS scans through the entire design to make sure that no two terminals have more than one path between them. Such checking assures that SPARCS will not run into a loop at run-time, but also introduces a tremendous number of unnecessary I/Os.
- (3) Most of the access patterns are predictable. Certain objects, such as the net instance in Figure 3.1, are accessed several times during navigation. Once this access pattern is known, the buffer manager can increase its containing page priority to increase the buffer hit ratio.

4. Simulation Model

A simulation model allows us to evaluate the performance impact of different clustering and buffering algorithms on different applications' characteristics. It also helps us answer certain what-if questions rapidly; for example, what is the relationship between the choice of the clustering algorithm and frequency parameters such as inheritance density and read/write ratio?

We used the Performance Analysis's Workbench System (PAWS) to construct the simulation because PAWS supports a number of high level primitives; for example, various queuing disciplines and a set of detailed statistics outputs. PAWS also allows us to refine and enhance the model easily. In the following section, we discuss the workload and our model in terms of PAWS primitives. We summarize all the modeling parameters in Section 4.2.

4.1. The Engineering DB Model

We represent our model of engineering database interaction in the PAWS language which allows the user to simply declare the characteristics of the system being modeled instead of coding detailed simulation algorithms.

The simulation model consists of several interacting model blocks. Transactions representing user requests flow between these model blocks carrying information about the work units. The major blocks, shown in Figure 4.1, are:

- (1) **Workstation Cluster:** a set of workstations representing interactive users and think times
- (2) **Workload Definition:** workload characteristics
- (3) **I/O subsystem:** a set of related components describing the I/O configuration
- (4) **Buffer Manager:** the buffer management module
- (5) **Cluster Manager:** various clustering algorithms management modules
- (6) **CPU:** the processor

Figure 4.1 shows the relationship between these blocks. A transaction starts at a *workstation cluster* node and submits a request to the file server after a predefined think time. The request is defined in the *workload definition* node. If the request is a read operation, the buffer manager searches through the buffer pool and issues a physical I/O if needed. The buffer searching phase flushes out some dirty pages as well as some transaction log records. Therefore, in the worst case, a logical I/O, generated in the workload definition node, can be translated into zero to 3 physical I/Os; one I/O flushes the dirty page, one I/O logs the transaction, and one I/O brings in the data. More detailed model description is in [CHAN89].

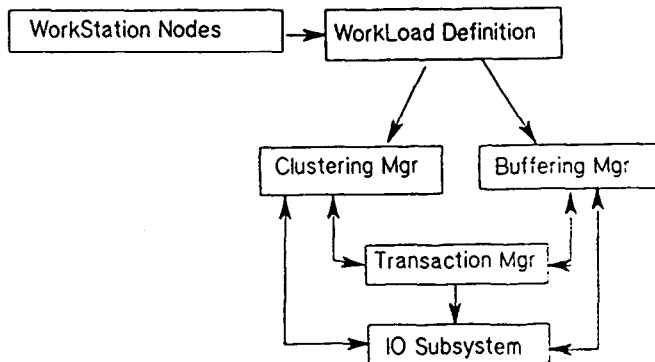


Figure 4.1 -- Simulation Model Overview

A few notes about our simulation model:

- (1) We model not only the buffer pool activities but also the transaction logging details to obtain realistic log I/O rates. A log record is constructed based on the size of the newly created or modified object. A circular in-memory log buffer is used and log records are flushed when the circular log buffer is full.
- (2) Object information such as object size, structural links with other objects, and containing page identifier, is maintained explicitly, allowing us to construct a realistic sample database used by all the buffering and clustering algorithms.
- (3) Actual physical I/O activities are modeled through the *I/O subsystem* node. The I/O path length and the disk service time are also modeled by the system.

The fundamental unit of recovery and concurrency control is the object and composite object. An object can be read by specifying an unique name or navigating along some structural relationships from other objects. If the target object for a read operation is composite, several logical I/Os may occur to retrieve the complete object. On the other hand, the write operation would only create or update a single object. Based on our measurements, we decided that every tool invocation (user session) is composed of 100 to 10000 queries with various read/write ratios.

Generally speaking, engineering design applications' procedure calls can be categorized into seven query types which are assigned to transactions in the workload definition phase: (1) Simple object lookup, (2) Component object retrieval, (3) Composite object retrieval, (4) Descendant version retrieval, (5) Ancestor version retrieval, (6) Corresponding objects retrieval, and (7) object insertion/deletion/updating.

In our simulation model, the checkin and checkout operations are modeled by these seven query types. For instance, a checkout operation may consist of several component object retrievals and one corresponding object retrieval. Similarly, a checkin operation invokes several object insertions and updating.

4.2. Parameters

Table 4.1 shows all the parameters modeled by the simulation model and their operating levels. They are divided into static parameters, which are fixed for all the simulation runs, and control parameters, which have various operating levels.

The operating levels and default values of these parameters are partially influenced by the structure density and read/write ratio observed in Section 3. The number of users sharing the same database is assumed to be 10 since most of the CAD/CAM designs can be partitioned into small units shared by less than 10 engineers. We used 10 disks in the server and assumed the page size to be 4 Kbytes. The database size is 500 Mbytes which is quite consistent with the specification in Sun engineering application benchmark [CATT88]. The 4 seconds think time for each user we believe represents what is the average think time in interactive computer-aided design environments.

We studied eight control parameters to understand their impacts on the overall system response time. Out of these eight control parameters, (F) and (G) determine the workload characteristics, (H), (I) and (J) control the clustering policy, and (K), (L), and (M) define the buffering strategy. We based the operating levels of Structure Density on actual observation from OCT tools access patterns. "Low-3" means that every structural retrieval returns fewer than or equal to three component or composite objects. "Med-5" means that more than 3 but fewer than 10 objects are returned through structural retrieval, and "high-10" means that 10 or more objects are

returned.

Label	Static Parameters	Default Value
A	Database Size	500 Mbytes
B	Page Size	4 Kbytes
C	Number of Users	10
D	Number of Disks	10
E	Think Time	4 seconds
	Control Parameters	Operating Levels
F	Structure Density	low-3,med-5,high-10
G	Read-write Ratio	5,10,100
H	Clustering Policy	No_Cluster, C_within_Buffer 2_IO_limit,10_IO_limit, No_limit
I	Page Splitting Policy	No_Split, Linear_Split, NP_Split
J	User Hint Policy	No_hint, User_hint
K	Buffering Policy	LRU, Context-sensitive,Random
L	Buffer Pool Size	100,1000,10000
M	Prefetch Policy	No_prefetch, P_within_buffer_pool P_within_Database

Table 4.1: Simulation Parameters

5. Simulation Results

The control parameters listed in Table 4.1 may interfere with each other. For instance, a higher structure density means more I/Os for the composite object retrieval and more candidate pages to consider during the object placement phase. A higher read/write ratio means that the extra I/Os caused by the writers during the clustering phase may be amortized by the readers and, consequently, that the overall system response time improves.

In Section 5.1, we discuss the run-time clustering effects using different clustering policies and various buffering strategies are studied in Section 5.2.

5.1. Run-Time Clustering Effect

The operating levels of buffering control parameters used in this clustering effects study are: No prefetch, 1000 buffers, and LRU buffer replacement policy. We first discuss the clustering effects when selected candidate page for object placement never overflow. That is, when the preferred candidate page is full, the storage manager chooses the next best candidate page which has enough space. Then, we study how various page overflow handling mechanisms affect clustering.

5.1.1. No Page Overflow

Figure 5.1 shows how various clustering policies affect system response time under different workloads. Some key observations are:

- (1) Run-time clustering always improves overall system response time. When both the structure density and read/write ratio are high, the response time is improved by 200%.
- (2) Setting I/O limits on potential candidate pages search is valid. When the structure density is low, *clustering with 2 I/O limitation* performs better or comparably to *clustering without I/O limitation*.
- (3) *Clustering within buffer pool* performs reasonably well when the structure density is low, but performance degrades to the *No_Clustering* case when the structure density is high. The buffer hit ratio also affects the performance of *clustering within buffer pool*. When the hit ratio is low, its performance is close to *No_Clustering*. However, performance improves

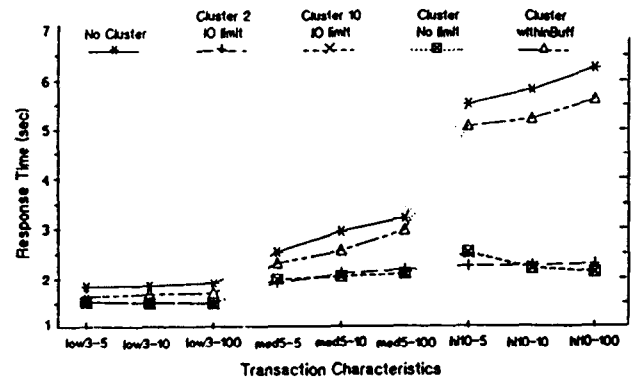


Figure 5.1 – Clustering Effects Analysis

when a better buffering policy is used (Refer to Section 5.2)

However, to determine the break-even point of the read/write ratio when *No_Clustering* has similar response time as any clustering mechanisms, we compare *No_Clustering* with *clustering without any I/O limitation*. The results are summarized in Table 5.1. Different structure densities have different break-even points, due to the amount of logical I/Os caused by writers during the clustering phase.

Structure Density	Read-write ratio
low-3	3.0
med-5	3.6
high-10	4.3

Table 5.1: Read-write ratio break-even points

Figures 5.2, 5.3 and 5.4 show the structure density effect on response time when the read/write ratio is fixed. In all cases, the response time rises sharply between medium structure density and high structure density when no clustering is done. Since a composite object retrieval may trigger 10 or more logical I/Os and the

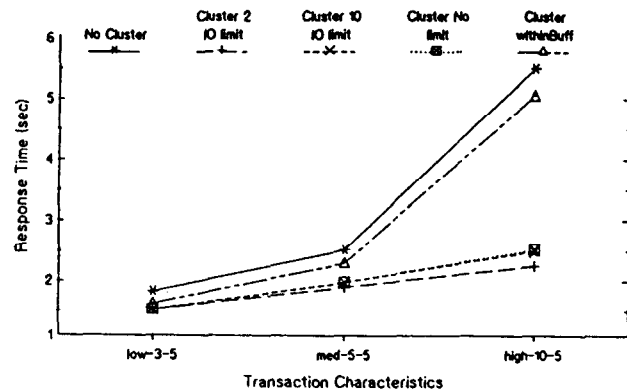


Figure 5.2 -- Clustering Effect Under R/W ratio 5

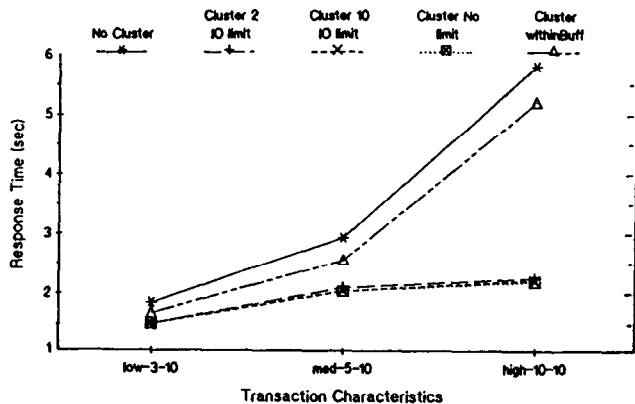


Figure 5.3 – Clustering Effect Under R/W ratio 10

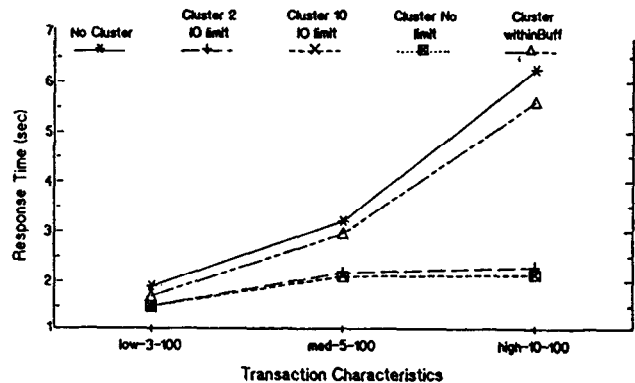


Figure 5.4 – Clustering Effect Under R/W ratio 100

possibility of buffer hits for these logical I/Os decreases, the increasing physical I/Os result in a higher response time. However, the response time rises slowly from a low structure density to a high structure density when any clustering mechanism except *clustering within buffer pool* is used. The low variation of response time over different structure densities is very critical when application access patterns are not fully predictable.

Clustering within buffer pool cannot perform well if the buffer pool hit ratio is low during the candidate page searching phase. By tracing the buffer pool usage, we found that the native LRU replacement policy frequently overlays the potential candidate page for new object placement and decreases the hit ratio. The buffer manager, if it understands the relationships among objects, may increase both the priority of these pages and the hit ratio. We will discuss the various buffering policies in Section 5.2.

When the read/write ratio is low, *clustering with a 2 I/Os limitation* provides the best response time in all structure densities as shown in Figure 5.2. For low structure density, *clustering with a 2 I/Os limitation* has the same response time as *clustering with no I/O limitation*. Notice that a low structure density implies fewer candidate pages to choose and fewer logical I/Os needed during clustering phase. As the structure density increases, the number of I/Os caused by searching candidate pages also increases. Such extra I/Os introduced in the clustering phase cannot be amortized by readers when the read/write ratio is low. Therefore, *clustering without I/O limitation* performs worse than *clustering with a 2 I/O limitation* in high structure densities.

Figure 5.3 shows that clustering with a limitation of 10 I/Os has the same response time as clustering with no I/O limitation for a medium structure density. Since the I/O limitation is larger than the maximum possible structure density, 10 I/Os behaves like no I/O limitation.

Clustering without I/O limitation performs consistently better than other clustering mechanisms when the read/write ratio is 100, as shown in Figure 5.4. Selecting a clustering mechanism based on the read/write ratio at run-time affords the best response time of both, setting either a small I/O limitation or no I/O limitation at all.

Clustering not only helps the retrieval but also reduces the number of I/Os in the transaction logging phase. This is shown in Figure 5.5 where we compare the number of transaction logging I/Os between no clustering and clustering without I/O limitation under different structure densities. When multiple updates occur on the same page within a transaction, the log manager needs to flush the original page only once. Since related objects are clustered on the same page, the probability of having multiple updates on the same page within a transaction increases, thus reducing the number of physical logging I/Os.

One interesting thing we learned from collecting OCT tools' access patterns is that the read/write ratio may vary across different phases of the same application. Therefore, we wanted to find out how the read/write ratio affects response time when the structure density is fixed.

Figure 5.6 shows that any clustering mechanism performs better than no clustering in the low structure density case. Clustering with and without I/O limitation perform similarly, and the variation of response time is very small. Therefore, for high read/write ratio applications having low structure density, *clustering with a 2 I/O limitation* may be the best choice. For applications having medium structure density, *clustering without I/O limitation* performs the best when the applications' read/write ratio is greater than 10. Notice that the response time of *clustering without I/O limitation* case changes little for all read/write ratios. Such a stable response time may be required by some real-time applications.

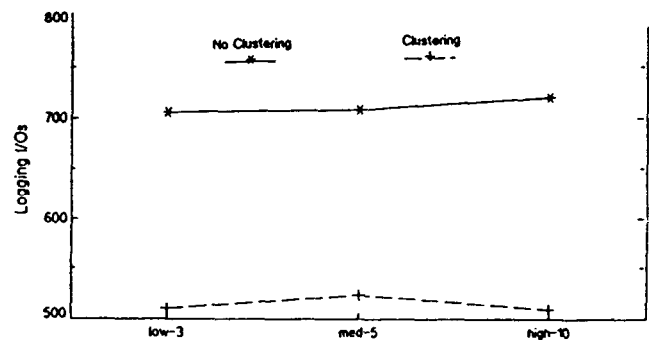


Figure 5.5 – Clustering Effect on Transaction I/Os




Figure 5.6 – Clustering Effect Under Low Structure Density

Figure 5.7 shows the clustering effect on response time when the structure density is high. Notice that the difference between the clustering within buffer case and the other clustering mechanisms becomes larger. This is due to the lack of potential candidate pages in the buffer pool, which reduces the effectiveness of run-time clustering.

5.1.2. With Page Overflow

When the chosen candidate page overflows with the newly inserted object, the system needs either to split the target page or to pick the next best candidate page. To split the target page, a page splitting algorithm is invoked, a new page is allocated, both the target page and the new page are modified, and the changes are logged by the system. Page splitting requires one extra I/O than no splitting to flush the newly allocated page later on and one extra log record that may trigger an I/O. Moreover, the page splitting algorithm needs extra CPU time, and the new page allocation may cause buffer pool contention.

As mentioned in Section 3, we have evaluated two different page splitting algorithms with the simulation model: *Linear Split* and *NP Split*. Figure 5.8 shows the effect of various page splitting algorithms when the clustering policy is clustering without I/O limitation. Solid line represents the response time for no page splitting case. Dotted line with plus sign is the response time when linear page splitting mechanism is applied. Dotted line with cross sign is the NP split case. When the read/write ratio is low, No splitting performs better than either splitting case. However, *Linear Split* provides the best response time when both the read/write ratio and the structure density are high. Both *NP Split* and *Linear Split* perform similarly when structure density is low because the number of arcs in the dependency graph is small, and the minor difference between the NP-complete solution and the linear solution is offset by other factors.

The objective of the page splitting algorithm is to minimize the expected access cost resulting from the page split. Ideally, the object partition having the minimum access cost should have the best response time. However, due to the extra I/Os and CPU time caused by the minimization process, the page splitting algorithm finding the minimum cost partition of objects may not perform best overall. Figure 5.9 shows the total cost difference from the Linear Split under different transaction characteristics. Notice that although the *NP Split* algorithm finds the minimum access cost partition of objects, its overall response time is worse than that of the *Linear Split* under high structure density.




Figure 5.7 – Clustering Effect Under High Structure Density

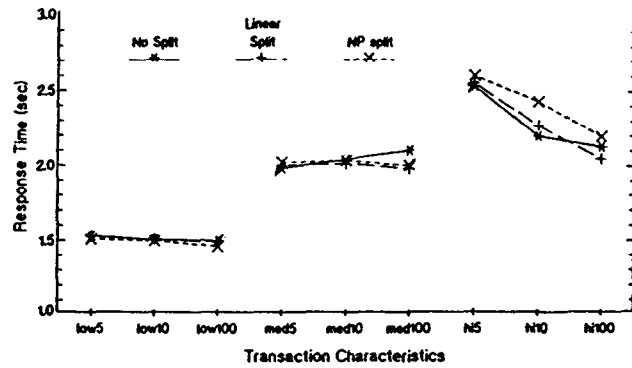


Figure 5.8 – Page Splitting Effects Analysis




Figure 5.9 – Cost Difference between NP Split and Linear Split

5.2. Impact of Smart Buffering

The operating levels of clustering control parameters used in this study of buffering effects are: clustering without I/O limitation and splitting page when the candidate page overflows. No user hints are provided for clustering and the number of buffers is 1000. The effect of buffer pool size on various buffering strategies are discussed in [CHAN89].

Three buffer replacement strategies, Context-sensitive, Random and LRU, are studied with three different prefetching strategies: prefetching within database, prefetching within buffer pool, and no prefetching.

Figure 5.10 shows how various buffering policies affect system response time under different workloads. Some key observations are:

- (1) The Context-sensitive buffer replacement policy always improves the overall system response time. When both the structure density and read/write ratio are high, *Context-sensitive with prefetching within database* outperforms *LRU with no prefetching* by 150% in response time.
- (2) Setting the scope on prefetching is a valid idea. When using prefetching within the buffer pool, the performance of the LRU and Random buffer replacement strategies is comparable to that of *Context-sensitive without any prefetching* in all workloads.
- (3) *Context-sensitive with prefetching within database* performs best, whereas *LRU with no prefetching* performs worst. A similar conclusion on the LRU buffer replacement strategy has been drawn in relational database systems.




Figure 5.10 – Buffering Effects Analysis

Six experiments are reported here: Context-sensitive buffer replacement policy with prefetching within database (C_p_DB), Context-sensitive with prefetching within buffer (C_p_buff), Random buffer replacement policy with prefetching within database (R_p_DB), Random buffer replacement policy with prefetching within buffer (R_p_buff), LRU buffer replacement policy with prefetching within database (LRU_p_DB), and LRU with no prefetching (LRU_no_p). Various transaction characteristics are evaluated. "hi100" means the transaction has high structure density and the read/write ratio is 100.

We have also examined the prefetching effect on response time when the buffer replacement algorithm is fixed. In all cases, *prefetch within database* performs best. This may not be intuitive since prefetching may bring in some pages not used by applications due to changing access patterns. However, because prefetching has made data available to applications in memory, the overall response time is improved. For object-oriented applications, paying extra I/Os to achieve better response time may be acceptable.

When the context-sensitive buffer replacement algorithm is used, as shown in Figure 5.11 *prefetch within buffer* has the same response time as *no prefetch* for transactions having low and medium structure densities. Notice that *prefetch within buffer* does not trigger any I/Os but only changes the buffer priority to reflect the future needs of data contained in these buffers. Since the context-sensitive buffer replacement algorithm would set up the appropriate priority based on knowledge of structure relationships and inheritance, *no prefetch* only causes very few buffer misses relative to *prefetch within buffer*. However, when the structure density becomes higher, *prefetch within buffer* better reflects the dynamics of the access pattern in how buffer priorities are set than *no prefetch* and has a better response time.

When the buffer replacement algorithm is not context-sensitive, prefetching becomes the only way to reflect the knowledge of structure relationships and inheritance in the buffer priority setting. Both *prefetch within database* and *prefetch within buffer* improve the buffer hit ratio and improve applications' response time. Although *prefetch within database* has more logical I/Os than *prefetch within buffer*, the overall response time is similar.

For more detailed analysis, refer to [CHAN89].

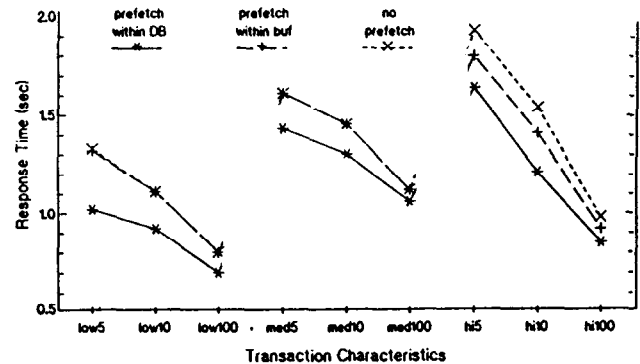


Figure 5.11 – Prefetching Effect under Context-sensitive Buffer Replacement Policy

6. Effect Analysis and Conclusions

Figure 6.1 shows the overall effect analysis for the eight control parameters specified in Table 4.1. There are 8! (i.e. 40,320) possible combinations, and for each one, we calculated how much, on the average, the response time has changed when the combined parameters go from low operating level to high operating level. Although the response time change can be positive or negative, only the absolute value is useful in this effect analysis. Larger absolute effect value means major effect on response time. When the effect value is zero, it implies that the corresponding control parameter(s) has minimum effect on the overall response time. Each blob represents the effects of a parameter or combined parameters on response time. For instance, "Structure density+buffering policy" represents a combined effect whereas "Read/write Ratio" is a single parameter effect. Most of the parameter or combined parameters have small response time change and are represented by a line of blobs centered in the figure. Two interesting observations are: 1) the structure density and buffering policy most influence the system response time, and 2) different page splitting algorithms have little influence on the response time.

Two control parameters, say A and B, are said to "interact" if the effect of A is *different* at different operating levels of B. They can be represented on an X-Y diagram where the X axis represents the different operating levels and the Y axis represents the effect (the response time). If two lines are parallel, there is no interaction between the corresponding control parameters. For intersected lines, we know there is a strong interaction between control parameters. If two lines do not intersect in the parameter range but are also not parallel, we say there is a minor interaction between the control parameters.




Figure 6.1 – Overall Effect Analysis

Figure 6.2 shows a subset of interaction analysis graphs. The key observations are:

- (1) There are no major interaction between any two factors. This means the control parameters we have chosen are quite independent.
- (2) There are minor interactions between: structure density and buffering policy, read/write ratio and clustering policy, read/write ratio and page splitting policy, structure density and clustering policy, structure density and page splitting policy, and page splitting policy and clustering policy.
- (3) There is no interaction between: buffering policy and clustering policy, buffering policy and page splitting policy, structure density and read/write ratio, and read/write ratio and buffering policy.

In this study, we have exploited inheritance and structural relationships for improved buffering and clustering. We have proposed a run-time clustering algorithm which under high structure density and large read/write ratio condition can improve system response time by a factor of 200%. We have also shown that it is effective to limit the amount of I/Os allowed to the clustering algorithm as it examines candidate pages for clustering at run-time. We studied the prefetching effect on response time under various buffer replacement policies and transaction characteristics. Context-sensitive buffer replacement policy with *prefetch within database* performs best whereas *LRU with no prefetching* performs worst.

We have also studied the effectiveness of *user hints* and how varying read-write ratios within a transaction affect the response time under various clustering and buffering algorithms [CHAN89]. However, the dynamics of access pattern of object-oriented applications are still not well understood. Furthermore, the physical representation of structural relationships and inheritance links among objects needs to be addressed.

Based on the simulation results shown in this study, we strongly recommend that future object-oriented DBMSs to (1) model structural relationships and inheritance links as first class objects, and (2) use the proposed dynamic clustering algorithm when the read-write ratio is high.

Acknowledgements

Our thanks to the SIGMOD referees for their helpful comments and suggestions on improving this paper. This research is partially supported by NSF under grant MIP 8706002.

References

- [ATWO85] Atwood, T., "An Object Oriented DBMS for Design Support Applications," Proc. IEEE COMPINT 85, Montreal, Canada, (Sept. 1985).
- [BATO85b] Batory, D., W. Kim, "Supporting Versions of VLSI CAD Objects," M.C.C. Technical Report, Austin, TX, (1985).
- [CHAN87a] Chang, E. E., R. H. Katz, "Inheritance in Computer-Aided Design Databases: Semantics and Implementation Issues," UCB Technical Report 87/377,(1987), Submitted for publication in *CAD Journal*.
- [CHAN87b] Chang, E. E., D. Gedyce, R. H. Katz, "The Design and Implementation of a Version Server for Computer-Aided Design Databases," *Software Practice and Experience*, in press.
- [CHAN89] Chang, E.E., "Effective Buffering and Clustering in Object-oriented Database Management Systems", Ph.D. Dissertation, May 1989.
- [HARR86] Harrison, David, P. Moore, R. Spickelmier, A.R. Newton, "Database Management and Graphics Editing in the Berkeley Design Environment", Proc. IEEE ICCAD, November 1986.
- [KATZ87] Katz, R. H., E. Chang, "Managing Change in a Computer-Aided Design Database," 13th Very Large Database Conference, Brighton, England, (Sept. 1987).
- [KATZ86a] Katz, R. H., E. Chang, R. Bhateja, "Version Modeling Concepts for Computer-Aided Design Database," ACM SIGMOD Conf., Washington, DC, (May 1986).
- [KATZ86b] Katz, R. H., E. Chang, M. Anwarrudin, "A Version Server for Computer-Aided Design Databases," ACM/IEEE 24th Design Automation Conf., Las Vegas, NV, (June 1986).
- [KIM87] Kim, Won, Jay Banerjee, Hong-Tai Chou, Jorge F. Garza, Darrel Woelk, "Composite Object Support in an Object-Oriented Database System," Proc. OOPSLA'87 Conf., Orlando, FL, (Oct. 1987).
- [MAIE86] David Maier, Jacob Stein, Allen Otis, Alan Purdy, "Development of an Object-oriented DBMS," Technical Report CS/E-86-005, (April, 1986).
- [PAWS83] Information Research Associates, Austin Texas, 1983, "Performance Analyst's Workbench System (PAWS)".
- [ZDON84] Zdonik, S. B., "Object Management System Concepts," Proc. 2nd SIGOA Conf. on Office Information Systems, Toronto, Canada, (June 1984).

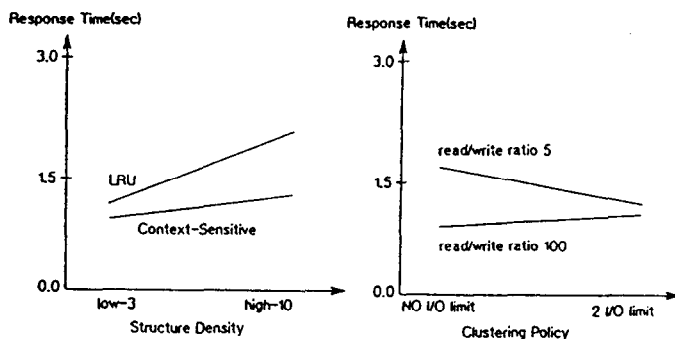


Figure 6.2 - Interaction Analysis Graph