

VDBMS: A Testbed Facility for Research in Video Database Benchmarking

Walid G. Aref, Ann C. Catlin, Ahmed K. Elmagarmid, Jianping Fan, Moustafa A. Hammad, Ihab F. Ilyas, Mirette Marzouk, Sunil Prabhakar, Yi-Cheng Tu, Xingquan Zhu *

Department of Computer Sciences, Purdue University, West Lafayette, IN. 47907-2066, USA.

Received: date / Revised version: date

Abstract Real world video-based applications require database technology that is capable of storing digital video in the form of video databases and providing content-based video search and retrieval. Methods for handling traditional data storage, query, search, retrieval and presentation cannot be extended to provide this functionality. The VDBMS research initiative is motivated by the requirements of video-based applications to search and retrieve portions of video data based on content, and by the need for testbed facilities to facilitate research in the area of video database management. In this paper we describe the VDBMS video database research platform, a system that supports comprehensive and efficient database management for digital video. Our fundamental concept is to provide a full range of functionality for video as a well-defined abstract database data type, with its own description, parameters and applicable methods. Research problems that are addressed by VDBMS to support the handling of video data include MPEG-7 standard multimedia content representation, algorithms for image-based shot detection, image processing techniques for extracting low-level visual features, a high-dimensional indexing technique to access the high-dimensional feature vectors extracted by image pre-processing, multimedia query processing and optimization, new query operators, real-time stream management, a search-based buffer management policy, and an access control model for selective, content-based access to streaming video. VDBMS also provides an environment for testing the correctness and scope of new video processing techniques, measuring the performance of algorithms in a standardized way, and comparing the performance of different implementations of an algorithm or component. We are currently developing video component wrappers

with well-defined interfaces to facilitate the modification or replacement of video processing components. The ultimate goal of the VDBMS project is a flexible, extensible framework that can be used by the research community for developing, testing and benchmarking video database technologies.

Keywords: Video database management system - Stream manager - Rank join - MPEG-7

1 Introduction

A significant and ever increasing portion of the information created today has audio-visual components, and most of it is now available in digital form. Real world video-based applications require database technology that is capable of storing this information in the form of video databases and providing content-based video search and retrieval. Methods for handling traditional data storage and retrieval cannot be extended to provide this functionality for video. Two standard approaches have been developed for handling video data: storage in video servers and storage as binary large objects. Storing video data in video servers with few data management capabilities has been widely adopted by video on demand (VOD) systems. The Fellini VOD system [29] supports real-time and non-real-time storage and retrieval of continuous media data (video, audio.) It focuses primarily on media streaming, and provides no support for content search and retrieval of the stored media. Many research VOD systems focus on efficient buffer management of raw video streams [8,23], and provide either very limited support or no support at all for content search and retrieval. It is not possible for such systems to address the integration of content search, retrieval and streaming. Integration of the search and streaming of video data is, however, supported by some commercial systems. Virage [38] provides textual extraction, encoding,

* This work was supported in part by the National Science Foundation under Grants IIS-0093116, EIA-9972883, IIS-9974255, IIS-0209120 and by the NAVSEA/Naval Surface Warfare Center, Crane.

search and streaming of continuous media. As a specialized video server, Virage offers a great deal of flexibility for applications that access video by textual content. On the other hand, it is difficult to integrate their video functionality with other types of data (e.g., relational data or high-dimensional data) or to express and execute complex declarative queries (e.g., SQL queries) on the stored media.

The second approach views video as a binary large object (BLOB), whose content is hidden from the system and for which no meaningful processing or optimization can be performed. BLOBs are supported by many industrial strength database systems. The problem with the BLOB representation is that once it leaves the database, it is handled by a user application in an application-specific manner, resulting in data type mismatches between the database and the application. Moreover, much important functionality, such as online and customized video views, query by content, similarity search queries, video editing functionalities and data abstraction cannot be provided as an integral part of a BLOB-based system. Some recent video research has focused on new applications of data management functionality to video data, such as mining videos [31] [40] and monitoring/tracking objects in video streams using continuous queries [15]. Such features will benefit from the efficient storage, query and indexing capabilities of a video database system which supports video data as a first class database object. This functionality cannot be satisfied by handling video data as BLOBs. The requirements must be addressed by building upon existing database technologies and extending them as needed to efficiently support video database functionality.

VDBMS provides a full range of functionality for video as a well-defined data type, with its own description, parameters and applicable methods. The development and integration of a video data type into the database management system achieves a clear separation between the video processing and database components. This allows video-based application design to focus on details of the application itself, while relying on the underlying video framework components for storage, search, retrieval, analysis and presentation of the video data. Video applications thus inherit all the powerful functionality generally provided by database management systems, including query processing, optimization, concurrency and recovery. Furthermore, VDBMS is a general purpose data management system, and supports extensions to include new data processing functionalities such as video mining and continuous queries.

VDBMS system components include a video pre-processing toolkit, a high-dimensional index manager, a stream manager, and a search-based buffer management policy. These VDBMS system components are described in this paper, and details can be found in the literature [2] [11] [14] [18]. We present VDBMS as a research platform because it provides an open and flex-

ible environment for investigating new research areas related to video database management, including the implementation, integration and evaluation of new and existing algorithms. Research problems that were addressed within the VDBMS environment to support the handling of video data include MPEG-7 document compliance for importing and exporting video features [1], algorithms for image-based shot detection [20], image processing techniques for extracting low-level visual features [11], camera motion detection algorithms [39], hierarchical video summarization strategies for abstracting video content, a high-dimensional indexing technique to access the high dimensional feature vectors extracted by image pre-processing, new multi-feature rank-join query operators for image similarity matching [18], a new real-time stream manager to admit, schedule, monitor and serve concurrent video stream requests periodically, an enhanced buffer management policy that integrates knowledge from the query processor to improve streaming performance [14], and an access control model that provides selective, content-based access to streaming video data [6].

The development of the VDBMS video database management research platform was motivated by the requirements of video-based applications to retrieve portions of video data based on content and by the need for testbed facilities to facilitate research in the area of video database management. While investigating, developing, and testing the fundamental components required to support full video database functionality, we utilized VDBMS as a testbed for integrating and evaluating video processing technologies from other sources. As such, the system has provided us with an environment for testing the correctness and scope of algorithms, *measuring* the performance of algorithms in a standardized way, and *comparing* the performance of different implementations of a component. The next step in VDBMS system development is the construction of video component wrappers with well-defined interfaces that allow video components to be easily modified or replaced. We also plan to provide the corresponding semi-automatic mechanisms for integrating these components into VDBMS. The ultimate goal of the VDBMS project is a flexible, extensible framework that can be used by the research community for developing, testing and benchmarking video database technologies.

We describe VDBMS system components in Sections 2 and 3. To demonstrate the usefulness of VDBMS as a testbed for video database benchmarking, Section 4 presents experimental studies and analysis for alternative techniques implemented within the VDBMS environment.

2 The Query Interface

A VDBMS query interface client supports end-user content-based query, search, retrieval and real-time

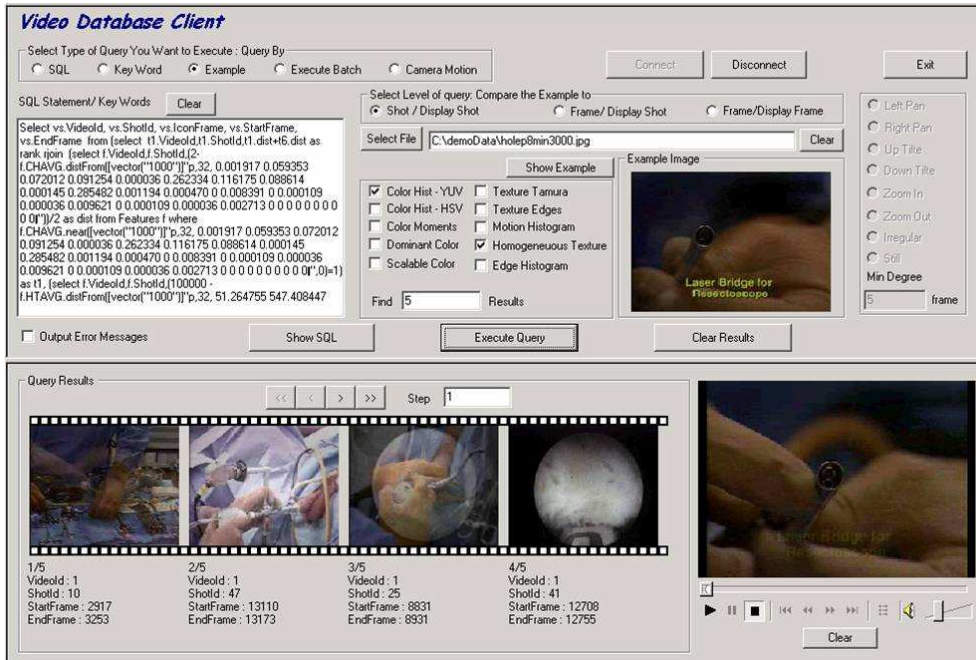


Fig. 1 VDBMS query interface.

streaming for the VDBMS video database server. End-users can query by image, query by camera motion type, query by keywords or specify an SQL statement. In image-based queries, users present an example image and query the database for images or shots "most similar" to the example based on any number and combination of the listed visual features. The features of the user's query image are extracted online and sent to the server for execution. Results can be either frame level (video frame images with similar visual features) or shot level (video shots with similar aggregate visual features, where feature aggregation is computed across shot frames.) Keyword queries are matched against video annotation data associated with logical video scenes. The query interface generates and displays the SQL equivalent of the user's query for all non-SQL queries. The VDBMS query processor returns a ranked list of results, where the user determines the number of top-ranked results to retrieve. Users can navigate an image skim of the results using any step size. When the user requests shot-level results, a key frame representing shot content is returned to the user, and the user can select the key frame to stream the shot directly from the database to the query interface media player.

Users access the VDBMS query interface using the Windows-based client shown in Figure 1. The client connects to the VDBMS system which resides on a Sun Enterprise 450 machine with 4 UltraSparc II processors, running the Solaris 5.6 operating system. The system memory is one Gbyte, with RAID disks of 170 Gbytes of storage. VDBMS functionality has been tested against more than 500 hours of medical videos obtained from the Indiana University School of Medicine. The medi-

cal videos are digitized, compressed into MPEG-1 format, processed off-line by the VDBMS pre-processing toolkit to generate image and content-based meta-data, and then stored together with their meta-data in the VDBMS database.

3 The Video Database Management System

The VDBMS database management system is built on top of an open source system consisting of Shore [36], the storage manager developed at the University of Wisconsin, and Predator [33], the object relational database manager from Cornell University. The VDBMS research group has developed the extensions and adaptations needed to support full database functionality for the video as a fundamental, abstract database data type. Key database extensions include high-dimensional indexing, video store and search operations, new video query types, real-time video streaming, search-based buffer management policies for continuous streaming, and support for extended storage hierarchies including tertiary storage. These extensions required major changes in many traditional database system components. Figure 2 illustrates our layered system architecture with its functional components and their interactions. The system consists of the object storage system layer at the bottom, the object relational database management layer in the middle, and the user interface layer at the top.

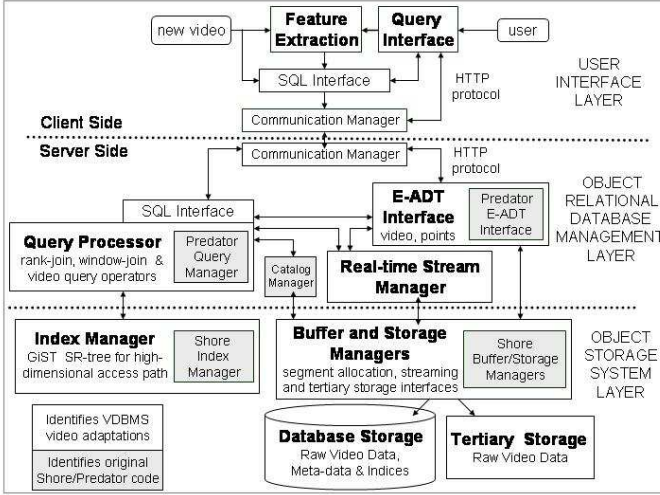


Fig. 2 VDBMS layered system architecture.

3.1 A Video Pre-processing Toolkit

The VDBMS video-preprocessing toolkit applies image and semantic processing to partition raw video streams into shots then associates the shots with extracted visual and semantic descriptors that represent and index the video content for searching. Preprocessing algorithms detect the video scene boundaries that partition the video into meaningful shots using a process that computes color histogram differences and incorporates a mechanism for dynamic threshold determination. Video shots are then processed to extract MPEG-7 compatible low-level visual feature descriptors, camera motion classification, spatial and temporal segmentation, representative key frames, and the semantic annotations of domain experts. The video, its features and indices are stored in the VDBMS database, along with physical metadata such as resolution for quality-of-service presentation. Our system follows the recent trend of representing the video content description in an XML-like format according to MPEG-7 [19] multimedia content descriptors. MPEG-7 is the worldwide standard for video content description, and has been incorporated as an integral part of VDBMS feature representation. VDBMS video pre-processing extracts nearly all low-level features defined by MPEG-7, including color histogram in both HSV and YUV formats, texture tamura, texture edges, color moment and layout, motion and edge histograms, dominant and scalable color, and homogeneous texture.

We are currently developing a wrapper that abstracts the extraction, representation, and query of features. This plug-in component allows users to define a new feature, supply its extraction (image processing) algorithm, and query against the feature for image similarity matching. Our wrapper and integration mechanisms incorporate the feature into the query interface, create the schema for database representation and apply the user-provided algorithm during video pre-processing and

image-based queries. We provide users with a graphical interface for defining and integrating video segmentation algorithms, feature extraction algorithms, camera motion classification techniques, and other video processing techniques that can be used for content representation and content-based retrieval. This will allow researchers to compare and evaluate alternate methods, improve existing algorithms or develop new ones.

3.2 High-dimensional Video Indexing

Since high-dimensional feature data is collected for each video frame and aggregated for each video shot, the meta-data that represents and indexes video content occupies more disk space than the video itself. The magnitude of this meta-data and its storage in the database as high-dimensional vectors present serious indexing and searching difficulties in the execution and optimization of feature-based queries. The VDBMS research group extended the indexing capability of Shore by incorporating the GiST v2.0 implementation [16] [17] [37] of the SR-tree as the high-dimensional index [5] [21] and modified the query-processing layer of Predator to access the Shore/GiST index. VDBMS added the vector ADT to be used by all feature fields, and implemented

```
CREATE GSR INDEX <table>_<fieldname> <table>
```

to create an instance of the GiST SR-tree for each field to be used as the access path in feature matching queries. The multi-dimensional indexing structure handles the high-dimensional feature vectors that are produced by visual feature extraction and used in image similarity searches.

We are currently building an interface to support a plug-in component for indexing techniques, so that alternative indexing mechanisms can be implemented, tested, and compared within the VDBMS system.

3.3 The Query Processor and Video Query Operators

The query processor was modified extensively to handle the new high-dimensional indexing scheme, as well as to support new video query operators and their integration into the query execution plan. VDBMS query processing must take into account the video methods and operators in generating, optimizing and executing query plans. Image similarity search is performed by issuing nearest neighbor queries to the high-dimensional access path.

In multi-feature image similarity queries, users generally present a sample image and query the database for images "most similar" to the example based on some collection of visual features. Results should be determined according to a combined similarity order [13] [26]. We have developed a practical, binary, pipelined query operator, NRA-RJ, which determines an output global

ranking from the input ranked video streams based on a score function [18]. Our algorithm extends Fagin’s optimal aggregate ranking algorithm, NRA, [10] by allowing the input to have a range of scores instead of exact scores. The output of NRA-RJ thus serves as valid input to other operators in the query pipeline, supporting a hierarchy of join operations and integrating easily into the query processing engine of any database system.

We created a new VDBMS query operator that encapsulates the rank-join algorithm in its *GetNext* operation. Each call to *GetNext* returns the next top element from the ranked inputs. The internal state information needed by the operator consists of a priority queue of objects encountered thus far, sorted on worst score in descending order. *GetNext* is binary (although this restriction is merely practical) and the algorithm holds for more than two inputs. Our modifications to the original NRA algorithm are the following:

- The right input list is a source stream that provides the operator with the ranked objects and their exact scores. The left input may not be a source list since it can be the output of another NRA-RJ operator. In this case, the score is expressed as a range, from worst to best. This means that *GetNext* must be able to handle a score range rather than an exact score from the left iterator.
- The parameter *k*, the number of requested output objects, is not known in advance, rather it increases for each call to *GetNext*. The modified algorithm first checks if another object can be reported from the priority queue without violating the stopping condition, and if not, moves deeper into the input streams to retrieve more objects.
- In each call to *GetNext*, the current depth of the caller is passed to the operator. This extra information assures synchronization among the pipeline of NRA-RJ operators.

The incremental and pipelining properties of our aggregation algorithm are essential for practical use in real-world database engines, and our new operator will help in implementing this type of join in ordinary query plans.

A modular interface for the integration of query operators into the VDBMS query processor is currently underway. The interface will support the integration of user-developed operators into the query execution plan, and will also support the performance evaluation and comparison of alternative algorithms for implementing query operators by allowing developers to identify performance metrics and test point locations for collecting measurements and statistics. In Section 4.2, we demonstrate this concept in the context of performance analysis for different algorithms that implement the multi-feature ranking query operator. The experimental study investigates scalability as well as time and space complexity, and discusses performance trade-off issues

3.4 The Stream Manager

The VDBMS stream manager is responsible for handling the special needs of video streaming. Each request for video data needs to be streamed with a predetermined rate (MPEG-1 needs, on average, a 1.5 Mbps display rate.) Violating the rate of streaming by either increasing or decreasing the display rate may result in overflow at the client buffer or hiccups at the client side. To hide the latency associated with access to disk storage, the stream manager streams part of the data (a *segment*) while pre-fetching the next segment into the memory buffers.

Since many stream requests are serviced simultaneously by the manager, resources such as memory buffers and disk bandwidth must be divided among the streams. This is achieved by serving each stream request periodically, and serving additional concurrent streaming requests within that period. Due to limited memory and disk bandwidth, the manager can only serve a specific number of requests within a single period. To serve requests in real-time, the segment referenced next should be retrieved into the buffer before the end of the current period. We have implemented a real-time stream manager above the buffer manager layer in VDBMS [1], and its functionality is as follows: (1) admit a new stream request if the maximum number of concurrent streams has not been reached, otherwise delay the request and re-try when one of the current requests finishes, (2) schedule segment pre-fetching by sending requests to the buffer manager. Each page of the allocated segment is fixed in the buffer pool until the page is streamed, (3) send the segment to the client according to a predetermined streaming rate. The segments are processed page-wise, and each page is unfixed and returned to the buffer manager after streaming the content, and (4) communicate with the query manager to keep track of search results. The stream manager is implemented as multi-threaded modules, and has well defined interfaces with the query engine, the buffer manager, and the Extensible Abstract Data Type (E-ADT) interface. The stream manager operates by issuing requests to the buffer manager, guiding the underlying buffer management policies, communicating with the query processor, and sending streams to clients at a specific rate.

3.5 A Search-based Buffer Management Policy

Continuous-media servers that support content-based search and retrieval use a main memory buffer to store the requested media streams before sending them on to the user. Buffering policies for media streaming have been investigated in several studies. Chang and Garcia-Molina [8] introduce a memory-efficient pre-fetching schedule based on fixing the time-displacement between pre-fetching requests. A recent study [23] proposes dy-

dynamic buffer allocation for media streaming that minimizes the memory requirement for concurrent media streams. The work in [12] presents the basic functionalities of buffer management for delay-sensitive multimedia data, and in [30] Ozden et al. describe changes needed by database management systems to support multimedia data. Replacement policies for media streams have been studied for target applications such as VOD [9] [25] [28], which are designed for streaming purposes only. Brown et al. [7] propose a goal-oriented buffer allocation for different database workloads, where a target goal (average execution time, for instance) is designated for each workload. In the VDBMS project, we have investigated a new buffer management policy which addresses the relationship between the searching and streaming processes of video data.

Caching parts of media streams that may be referenced in the near future enhances streaming performance in two ways: it reduces the number of references to disk storage and it minimizes delay associated with the start of streaming. However, precise caching decisions are often difficult to make. Optimal pre-fetch and replacement policies would pre-fetch the data before its first reference and replace the data block that will not be referenced for the longest time [30]. An obvious difficulty is the policy's dependence on knowledge about expected streams, which is generally not available. In the case of video streaming, however, there is an inherent connection between query processing and streaming. Choices for streaming are usually based on query results, and this relationship can be used by the buffer manager to pre-fetch and cache pages expected for reference.

We have developed an efficient buffer management policy that uses feedback from the search engine to make more accurate replacement and pre-fetching decisions. Top-ranked query results from the query processor are used to predict future video streaming requests, and a weight function [3] determines candidates for caching. By integrating knowledge from the query and streaming components, VDBMS can achieve better caching of media streams, thus minimizing initial latency and reducing disk I/O. Many factors must be considered when basing pre-fetching or replacement decisions on search results. Streaming based on the search context is probabilistic: new streaming requests can be based on any of the search results or even on none of them. Furthermore and as a result of sharing the caching space by pages for current as well as expected streams; there will be increased overhead in the replacement policy associated with balancing the space assigned to each.

In our search-based replacement policy, pages in the buffer pool that are referenced by either current or expected streams are considered for caching. To maximize the number of caching pages, we replace the page that will be referenced by current or expected streams after a long period of time has passed. Also, we prefer caching pages that will be reference by current streams to those

that will be referenced by expected streams by assigning higher *keep* weight values to the current streams. Lookup tables contain pointers to expected streams, which are collected from the search results and checked by the stream manager for matches when determining pages to replace. With knowledge collected in lookup table for expected streams, we predict with high probability that one of the expected streams will be requested. The stream manager tracks the utilization of the streaming period, and utilizes any fraction of the streaming period unused by current streams to pre-fetch the first segment of the top ranked expected streams into the memory buffer. If an expected stream becomes an actual request, most of the pages in the first segment would already be cached in the buffer pool and as a result, the number of references to lower level storage would be significantly reduced. The pre-fetching policy does not introduce much overhead, since it operates only during idle period time, utilizing unused and reserved streaming resources.

The performance of the search-based policy was evaluated by investigating the effects of buffer management on the number of I/Os when referencing the first segment of a requested stream. Experimental results are presented in Section 4.1. They show that initial latency of the search-based policy is reduced on the average by 20% when compared with traditional policies.

3.6 Extended Storage Hierarchies

Video database storage and buffer managers handle huge volumes of data with real time constraints [25] [28]. In VDBMS, the buffer pools are divided between the database buffer area and the streaming area where requests for streams are serviced. Extended buffer management handles multiple page requests with segment allocation (instead of the traditional page-based approach) for the large streaming requests from the stream manager. An interface between the buffer manager and the stream manager is used to exchange information that guides buffer caching for stream requests. The storage manager was extended to perform necessary video operations and process both real-time and non-real time requests. VDBMS methods for handling extended storage hierarchies support transparent, real-time access to buffer, disk and tertiary storage. Different caching levels on buffer and disk storage enhance access for frequently referenced data, and a tertiary storage server manages access to tertiary resident data, making it directly accessible to the VDBMS system.

The tertiary storage and cache disk managers are implemented underneath the storage volume manager. A dedicated disk partition is used for caching hot items in tertiary storage, and the cache disk manager maintains and reports these items. The tertiary storage manager communicates with a PowerFile manager (remote NT process), which locates the requested page in its

changer, loads the CD/DVD, and sends the requested data through the local network via TCP. Due to the overhead for reading from the device, the basic transfer unit is a block instead of a page. The block is copied to the cache disk, and the first page is sent to the buffer pool. DVD jukeboxes can be daisy-chained, giving VDBMS access to terabytes of data.

4 A Testbed for Video Database Benchmarking

While investigating and implementing components to support full video database management, we have utilized VDBMS to investigate, integrate, validate, compare and evaluate alternate video processing techniques and technologies. To illustrate the effectiveness of the current VDBMS system for new component integration, validation, and performance evaluation, we briefly describe three recent research projects carried out within the VDBMS environment. The contribution of these and other experimental studies to the understanding of video processing within the database environment is the motivation for our effort to create a complete testbed facility for video database benchmarking.

4.1 Validation of a Buffer Management Policy

To validate the search-based buffer management policy in a heavy workload environment, we execute 32 simultaneous clients. Each client submits an image-based query to VDBMS and receives a collection of key frame representing the results of a shot-based image similarity search. The client delays for a random period (uniformly distributed between 10 to 20 seconds) after retrieving the results, and then submits a streaming request for one of them. We assume the client plays a shot selected from the four top-ranked results 80% of the time. The VDBMS stream manager admits the streaming request if possible; otherwise the request is delayed until one of the current streams has finished. The client immediately submits a new search request following the streaming of the selected shot, so that a heavy load situation is maintained. The search results are synthesized by random selection of 10 candidate shots from the database. The random selection provides an upper bound for the performance of our policy. Our *keep* weight is set to three if a page is referenced by an expected stream from the top ranked results, and four if the page is referenced by a current stream. Higher values for the keep parameter are not recommended since they lead to excessive looping over buffer pages to find replacement candidates. The experimental data consists of eight one-hour videos, compressed in MPEG-1 format with a total size of five Gbytes. Each video has been pre-processed into shots with lengths between 5 and 10 minutes. We set the page size to 8Kbytes, the segment size to 30 pages, and the

maximum number of concurrent streams to 16. Each experimental run lasts for 30 minutes, and the total number of buffer references is approximately 500,000.

We compare the performance of the following policies:

- Search-based replacement (SrchBR): pages are cached if referenced by current or expected-stream requests
- Search-based pre-fetching and replacement (SrchBPR): first segment of expected-stream is pre-fetched; pages are cached if referenced by current or expected-stream requests
- Stream-based replacement (StrmBR): pages are cached only if referenced by a concurrent stream request.
- Use&Toss: pages are candidates for replacement immediately after use.

Figure 3(a) shows the effect of the buffer policies on reducing the number of I/Os when referencing the first segment of the stream. For each first segment, we measure the percentage of pages found in the buffer as we increase the buffer size from 10 to 25 Mbytes. The figure shows that SrchBPR caches about 25% of the total pages of new streams based on the search results; that is, the initial latency is reduced by 25%. Although SrchBR achieves better results than StrmBR and Use&Toss, it caches only those pages either used by current streams or referenced by expected streams, and therefore the improvement is smaller than that of the pre-fetch policy. StrmBR has no knowledge of expected streams and performs about the same as Use&Toss. In Figure 3(b), the buffer size is fixed at 25 Mbytes, and we measure the reduction in I/O when referencing the first segment of the stream as the number of stored videos is increased from two to eight. SrchBPR achieves the best performance, as high as 40% reduction in the number of I/Os. This improvement results from both pre-fetching and replacement strategies, since more common data now exists between current and expected-streams. As the number of videos increases, the chance for interaction decreases. Thus the improvement is dominated by the positive effects of pre-fetching. The effect of the replacement policy is obvious in SrchBR and StrmBR, as both reduce the I/Os with small data sets. With larger data sets, StrmBR and Use&Toss contribute similarly to the reduction of I/O, since both have no knowledge about expected streams. The short duration of streamed segments represents an obstacle for replacement algorithms that depend only on current streams for two reasons: 1) in large data set with uniform access patterns, common pages are infrequent, and 2) common pages generally exist within a short interval of each other (intervals are bounded, on average, by half the length of a shot). Replacement policies based on caching common pages between current streams will thus have a small number of pages to recommend for caching.

Figure 3(c) shows the relative improvement in the buffer hit ratio for policies based on current streams.

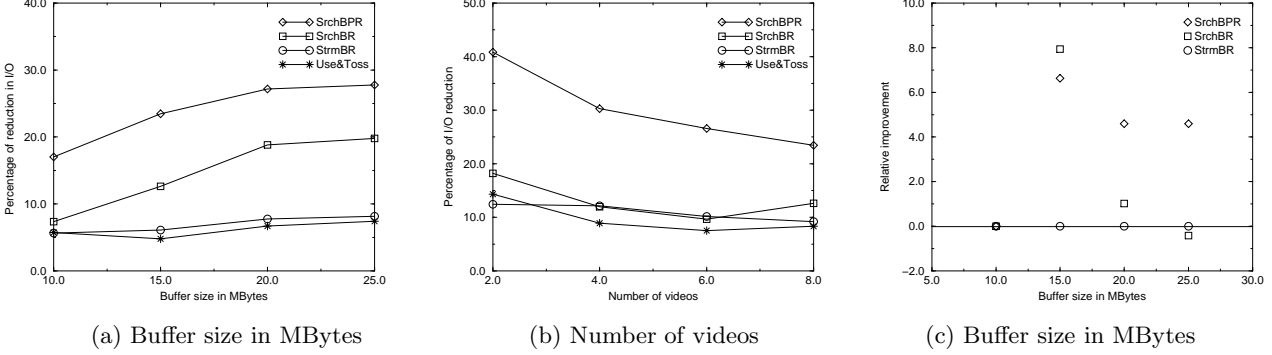


Fig. 3 Performance evaluation of buffer management policies: (a) reduction in I/O as buffer size changes, (b) reduction in I/O as number of videos changes, and (c) relative improvement in buffer hit ratio as buffer size changes.

As the buffer size increases, more space is available to cache the data and the chance of replacement is decreased. With small buffer sizes, pages are replaced more frequently and the improvement achieved with search-based policies such as SrchBPR and StrmBR becomes significant.

4.2 Performance Evaluation of Rank-Join Query Operators

We implemented three state-of-the-art rank-join algorithms as query operators in VDBMS for an extensive empirical study to evaluate operator performance and trade-off issues in executing multi-feature queries. Our experimental study compares the NRA-RJ operator developed by the VDBMS research group [18], the J* operator introduced by Natsev et al. [27], and (for a baseline comparison) the non-pipelined version of the NRA algorithm as a multi-way rank-join operator, MW-RJ [10]. Although most query optimizers are restricted to binary operators, MW-RJ provides a reference line for the best possible performance. We investigated scalability as well as time and space complexity between the algorithms for executing a join of multiple ranked inputs (any number and combination of features) on the stored video objects.

The following multi-feature query for the k top-ranked results was issued against the VDBMS features:

Retrieve the top k video shots "most similar" to a given image based on m visual features

The query evaluation plan has m nearest neighbor (NN) operators on m different visual features, and $m - 1$ rank-join binary operators are used, where the results of one operator are pipelined to the next operator in the pipeline. The number of features m in our study varies from 2 to 6, and the number of top-ranked results k varies from 5 to 100. To evaluate the operators, we used the following performance metrics: (1) query running time for retrieving the top matching k output results, (2) size of the buffer maintained by the operator, and (3) number of database accesses in disk pages. While

the number of database accesses should give a good indication of the time complexity of the operator, the experiments show a significant CPU time complexity difference between the two operators that affects the total running time, especially for small numbers of inputs. Figures 4 and 5 give performance comparisons for NRA-RJ, J* and MW-RJ, for $m = 2$ and $m = 3$, respectively, where m is the number of input sources that give a pipeline of length $m - 1$. For $m = 2$, NRA-RJ is identical to MW-RJ since there is no pipeline. Figure 4(a) compares the total running time of the NRA-RJ and J* operators. The J* algorithm has a significant CPU overhead due to the execution of its underlying A* graph search algorithm, which considers more join combinations. Thus, NRA-RJ shows a faster execution time. Both operators are nearly equal in the database access count depicted in Figure 4(c). NRA-RJ has a smaller maximum queue size than that of J*, as shown in Figure 4(b), and the difference increases as k increases (i.e., as more results are requested). The difference in the maximum queue size and in the execution time can be explained by the fact that the J* algorithm has to consider more join combinations than NRA-RJ since it was developed for a general join condition. When used in self-join problem settings, the generality of the J* algorithm causes expensive unnecessary computations that increase both the queue size and the running time.

Figure 5 compares the NRA-RJ, J*, and MW-RJ operators for $m = 3$. Figure 5(a) shows that NRA-RJ still outperforms J* in total running time, and the pipeline does not affect the speed of the NRA-RJ operator when compared with MW-RJ. For the maximum queue size given in Figure 5(b) and the number of database accesses given in Figure 5(c), we make the following observations:

- NRA-RJ has a larger maximum queue size and more database accesses than MW-RJ. To understand this difference, we clarify how NRA-RJ operates in a pipeline of 3 inputs. Figure 6 shows NRA-RJ with three input streams, L_1 , L_2 and L_3 . When the top NRA-RJ operator, OP_1 , is called to produce the next top ranked ob-




Fig. 4 Performance comparison for NRA-RJ and J* when $m=2$.

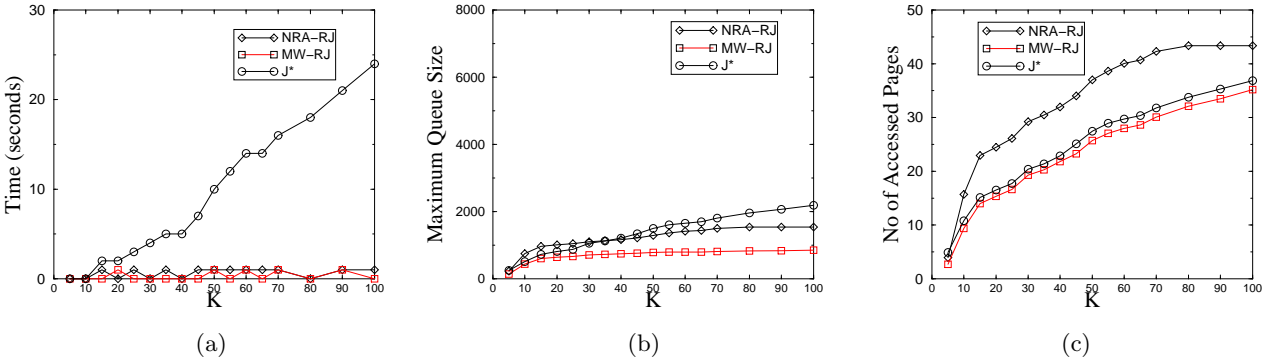


Fig. 5 Performance comparison for NRA-RJ, J* and MW-RJ for $m=3$.




Fig. 6 A query pipeline with $m=3$.

ject, several GetNext calls for the left and right children are invoked. According to NRA-RJ's *GetNext* algorithm, OP_1 gets the next tuple from its left and right children at each step. Hence, OP_2 will be required to deliver as many top ranked objects for L_2 and L_3 as for L_1 . These calls to the ranking algorithm in OP_2 force L_2 and L_3 to retrieve unnecessary objects, which results in larger queue sizes with more database accesses. We refer to this as the *local ranking* problem, that is, the NRA-RJ operator in the early pipeline stages tends to retrieve more database

objects in order to deliver as many ranked tuples as required by the next NRA-RJ operator.

- The J* operator has less database access cost than NRA-RJ, and close to the cost of MW-RJ, despite NRA-RJ's *local ranking* problem. In contrast to NRA-RJ, the J*'s algorithm does not retrieve equal numbers of objects from its left and right children.
- For the same reason that J* has less disk accesses than NRA-RJ, J* starts with smaller maximum queue size than NRA-RJ. However, as in the case for $m = 2$, J* begins to save many candidate join combinations in the queue, causing its maximum queue size to become larger than that of NRA-RJ as k increases. This also explains the fact that J* has a larger queue size than MW-RJ, even though both are retrieving almost the same number of database objects, as shown in Figure 5(c).

We now evaluate the scalability of the two pipelined operators with respect to the length of the query pipeline m . By fixing $k = 20$, the operators NRA-RJ and J* are again compared with respect to our three chosen performance metrics. As m increases from 2 to 6, NRA-RJ has a larger queue size because of the increased local ranking overhead in the pipeline. As NRA-RJ encounters greater database access, I/O cost begins to dominate total running time. The overhead finally affects the running time enough to make NRA-RJ performance worse than J*,




Fig. 7 The optimized NRA-RJ operator.

demonstrating clearly that J^* is scalable in terms of increased ranked inputs while NRA-RJ is not. Our evaluation of the performance of NRA-RJ led to an important insight: we must minimize the excessive local ranking calls in earlier stages of the pipeline. Our solution was to unbalance the depth step in the operator children. We changed the NRA-RJ GetNext algorithm to reduce the local ranking overhead by changing the way it retrieve tuples from its children, that is, to require less expensive GetNext calls to the left child, which is also an NRA-RJ operator. Using different depths in the input streams had a major effect on the performance. Figure 7 shows the comparison between the modified NRA-RJ, the J^* and the MW-RJ operator. In the optimized version of the NRA-RJ operator, one tuple is retrieved from the left NRA-RJ child for each p tuples retrieved from the right input child (in the figure, $p = 2$.) The optimized NRA-RJ operator showed significant performance improvements in both the maximum queue size and in the number of database accesses, due to the reduction of local ranking overhead in the inner pipeline stages. With this improvement, the optimized NRA-RJ operator is superior to the J^* operator, even for large m . The optimized NRA-RJ operator is an order of magnitude faster, has less space requirements, and has a comparable number of disk accesses.

4.3 A Component for MPEG-7 Document Compliance

The existing implementation for VDBMS feature extraction and database representation applies MPEG-7 to descriptors and description schemes, using an XML-like format to define the semantic and image-based information that identifies video content. The video pre-processing toolkit extracts nearly all low-level features defined by MPEG-7 as standard, and the VDBMS query interface allows to users to retrieve video shots based on any combination of these features. Queries combining multiple low-level features can be used to approximate high-level content-based searches.

To accommodate the representation and query of MPEG-7 features not currently extracted by VDBMS

video pre-processing, we have developed an XML wrapper that can import any MPEG-7 document specified with Data Definition Language (DDL) and map its descriptors to the VDBMS object-relational database schema. The wrapper also supports the export of VDBMS extracted features and other meta-data from the database as an MPEG-7 document. The XML wrapper enables the VDBMS system to make use of any available pre-extracted meta-data formatted as MPEG-7 documents without preprocessing the video itself. In addition, features that VDBMS video pre-processing does not extract (such as event-based and other semantic features) can be integrated, represented and queried as VDBMS meta-data via this mechanism. A document import function takes as input a user-supplied MPEG-7 document, which is generated using Multimedia Description Schemes (MMDS) and contains the high and low-level feature descriptors. The document is passed through the VDBMS MPEG-7 wrapper to extract, parse and map the descriptors to the VDBMS database feature schema. The video and its documented MPEG-7 features are then stored inside the database where they can be used for image and content-based queries. An export function extracts existing feature descriptors from the database and sends them through the wrapper where they are mapped to the MPEG-7 descriptors. The generated document can be used by other video processing tools or databases.

5 Conclusion

In this paper, we present a video database research initiative that resulted in the successful development of a video database management system which provides comprehensive and efficient capabilities for indexing, storing, querying, searching, and streaming video data. Our fundamental concept was to support a full range of functionality for video as a fundamental, well-defined abstract database data type. Research problems that were addressed by VDBMS to support the handling of video data include MPEG-7 standard multimedia content representation, algorithms for image-based shot detection,

image processing techniques for extracting low-level visual features, a high-dimensional indexing technique to access feature vectors extracted by image pre-processing, multimedia query processing and optimization, new query operators, a real-time stream manager, a search-based buffer management policy, and an access control model for selective, content-based access to streaming video data. We have also used VDBMS as a testbed for integrating and evaluating video processing techniques and components. As such, the system has provided us with an environment for testing the correctness and scope of algorithms, measuring the performance of algorithms in a standardized way, and comparing the performance of different implementations of components. The use of VDBMS as a testbed facility was illustrated by performance studies to investigate and analyze alternative implementations of video database processing methods.

We are currently constructing video component wrappers with well-defined interfaces to facilitate the modification or replacement of video processing components. We are also developing semi-automatic mechanisms for integrating these components into VDBMS. The ultimate goal of the VDBMS project is a flexible, extensible framework that can be used by the research community for developing, testing and benchmarking video database technologies.

References

1. Aref W, Catlin A, Elmagarmid A, Fan J, Hammad M, Ilyas I, Marzouk M, Zhu A (2002) A Video database management system for advancing video database research. In International Workshop on Management Information Systems. Tempe, Arizona, USA.
2. Aref W, Catlin A, Elmagarmid A, Fan J, Guo J, Hammad M, Ilyas I, Marzouk M, Prabhakar S, Rezgui A, Teoh S, Terzi E, Tu Y, Vakali A, Zhu X(2002) A distributed server for continuous media. In Proceedings of the 18th International Conference on Data Engineering. San Jose, California., USA.
3. Aref W, Kamel I, Ghandeharizadeh S (2001) Disk scheduling in video editing systems. *IEEE Trans. on Knowledge and Data Engineering*. 13(6). pp. 933-950.
4. Beckmann N, Kriegel H, Schneider R, Seeger B (1990) The R* -tree: an efficient robust access method for points and rectangles. *SIGMOD Record*, ACM Special Interest Group on Management of Data, 19(2): pp. 322-331.
5. Berchtold S, Bhm C, Jagadish H, Kriegel H-P, Sander J(2000) Independent quantization: An index compression technique for high-dimensional data spaces. In Proceedings of the 16th International Conference on Data Engineering. San Diego, California, USA. pp. 577-588.
6. Bertino E, Hammad M, Aref W, Elmagarmid A (2000) An access control model for video database systems. In Proceedings of the 9th International Conference on Information and Knowledge Management. pp. 336-343.
7. Brown K, Carey M, Livny M (1996) Goal-oriented buffer management revisited. In Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data. Montreal, Quebec, Canada.
8. Chang E, Garcia-Molina H (1997) Effective Memory Use in a Media Server. In Proceedings of the 23rd International Conference on Very Large Data Bases. Athens, Greece.
9. Dan A, Sitaram D (1996) A generalized interval caching policy for mixed interactive and long video environments. In the IS&T SPIE Multimedia Computing and Networking Conference. San Jose, CA, USA.
10. Fagin R, Lotem A, Naor M (2001) Optimal aggregation algorithms for middleware. In Proceedings of the 20th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems. Santa Barbara, California, USA.
11. Fan J, Aref W, Elmagarmid A, Hacid M-S, Marzouk M, Zhu, X (2001). Multiview: Multi-level video content representation and retrieval. *Journal of Electrical Imaging*, Vol. 10, No. 4. pp. 895-908.
12. Gemmell J, Christodoulakis S (1992) Principles of Delay Sensitive Multimedia Data Storage and Retrieval. In *ACM Transactions on Information Systems*, Vol(1) No(1) pages 51-90.
13. Guntzer U, Balke W-T, Kiessling W (2000) Optimizing multi-feature queries for image databases. In Proceedings of the 26th International Conference on Very Large Databases. September 10-14, 2000. Cairo, Egypt. p. 419-428.
14. Hammad M, Aref W, Elmagarmid A(2002) Search-based buffer management policies for streaming in continuous media. In Proceedings of the IEEE International Conference on Multimedia and Expo. Lausanne, Switzerland.
15. Hammad M, Aref W, and Elmagarmid A(2003). Stream window join: tracking moving objects in sensor network databases. In Proceedings of the 15th International Conference on Scientific and Statistical Database Management.
16. Hellerstein J, Naughton J, Pfeffer A(1995) Generalized search trees for database systems. In Proceedings of the 21st International Conference on Very Large Data Bases. Zurich, Switzerland.
17. Ilyas I, Aref W (2001) SP-GiST: An extensible database index for supporting space partitioning trees. *Journal of Intelligent Systems (JIIS)*. 17(2-3). pp. 215-235.
18. Ilyas I, Aref W, Elmagarmid A(2002). Joining ranked inputs in practice. In Proceedings of the 28th International Conference on Very Large Databases. Hong Kong, China.
19. ISO/IEC/JTC1/SC29/WG11: Text of ISO/IEC 15938-3 Multimedia Content Description Interface ? Part 3: Visual. Final Committee Draft. Document No. N4062. Singapore. March 2001
20. Jiang H, Helal A, Elmagarmid A, Joshi A(1998). Scene change detection for video database systems. *Journal on Multimedia Systems*, 6(2), pp.186-195.
21. Katayama N, Satoh S (1997). The SR-tree: An index structure for high dimensional nearest neighbor queries. *SIGMOD Record*, ACM Special Interest Group on Management of Data, 26(2).
22. Kumar P, Babu G (1998). Intelligent multimedia data: data + indices + inference. *Multimedia Systems*. 6(6). pp.395-407.
23. Lee L, Whang Y, Moon S, Song Y (2001). Dynamic Buffer Allocation in Video-on-Demand Systems. In Proceedings of the ACM SIGMOD International Conference on Management of Data, Santa Barbara, CA, USA.
24. Lin T, Zhang H (2000). Automatic Video Scene Extraction by Shot Grouping. In Proceedings of the ICPR.

25. Moser F, Kraiss A, Klas L (1995). A buffer management strategy for interactive continuous data flows in a multimedia dbms. In Proceedings of the 21st International Conference on Very Large Data Bases. Zurich, Switzerland. pp. 275-286.
26. Nepal S, Ramakrishna M (1999) Query processing issues in image (multimedia) databases. In Proceedings of the 15th International Conference on Data Engineering. Sydney, Australia. p. 22-29. IEEE Computer Society.
27. Natsev A, Chang Y-C, Smith J, Li C-S (2001). and Vitter, J. Supporting incremental join queries on ranked inputs. In VLDB'01 Proc. of 27th International Conf. on Very Large Data Bases. Rome, Italy.
28. Ozden B, Rastogi R, Silberschatz A (1996). Buffer replacement algorithms for multimedia storage systems. In Proc. of IEEE International Conf. on Multimedia Computing and Systems. pp. 172-180.
29. Ozden B, Biliris A, Rastogi R, Silberschatz A (1994) Fellini: A Low-cost Storage Server for Movie on Demand Databases. In Proceedings of the 20th International Conference on Very Large Databases.
30. Ozden B, Rastogi R, Silberschatz A (1997). Multimedia Support for Databases. In Proceedings of the 16th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems. Tucson, Arizona, USA.
31. Pan J-Y, Faloutsos C (2003) GeoPlot: spatial data mining on video libraries. In Proceedings of the International Conference on Information and Knowledge Management, McLean, VA, USA.
32. Rui Y, Huang T, Mehrotra S (1999) Constructing table-of-content for video. ACM MSJ. 7(5). pp. 359-368.
33. Seshadri P (1998) Predator: A resource for database research. SIGMOD Record. 27(1). pp. 16-20.
34. Smith J (1978) Sequentiality and prefetching in database systems. ACM Trans. on Database Systems. 3(3). pp. 223-247.
35. Stonebraker M (1981) Operating system support for database management. CACM. 24(7). pp. 412-418.
36. Storage Manager Architecture. Shore Documentation, Computer Sciences Department. UW-Madison, June 1999.
37. Thomas M, Carson C, Hellerstein J (2000) Creating a Customized Access Method for Blobworld, In Proceedings of the 16th International Conference on Data Engineering, San Diego, California, USA.
38. Virage. <http://www.virage.com/>.
39. Zhu X, Elmagarmid A, Xiangyang X, Catlin A. InsightVideo: Toward hierarchical content organization for efficient video browsing, summarization and retrieval. IEEE Transactions on Multimedia Journal. To appear.
40. Zhu X, Fan J, Aref W, Catlin A, Elmagarmid A (2003) Medical Video Mining for Efficient Database Indexing Management and Access. In Proceedings of the 19th International Conferences on Data Engineering. Banalore, India.