Very Large Databases

Minos N. Garofalakis
Bell Laboratories
minos@bell-labs.com

Renée J. Miller
University of Toronto
miller@cs.toronto.edu

A growing number of database applications require on-line access to very large volumes of data to perform a variety of tasks. For example, telecommunication service providers need to store Terabytes of phone call data to satisfy the requirements of billing, fraud detection, and strategic planning. Furthermore, many of these applications require support for new types of digital data, such as images and video. For example, a detailed requirements analysis for NASA's Earth Observing System indicates that at the turn of the century, the daily growth in stored image data will be 2.7 Terabytes, and the total stored volume will exceed 1.5 Petabytes. In this chapter, we identify such new application domains for database management and discuss the issues that arise from large data volumes, application-specific requirements, and new types of data.

Application-Oriented Database Management Systems

Database management systems (DBMS) are designed to provide the data storage and manipulation functions common to tasks that depend on very large volumes of data. Economic and technological changes, including the development of high speed networking, are fueling a new family of data-intensive applications. Traditional DBMS applications, such as banking applications, required fast access by multiple users to large, dynamic datasets. To meet these requirements, traditional DBMS support on-line transaction processing (OLTP), using transactions as the basic mechanisms for ensuring data consistency in the face of concurrent updates by a host of users. The data are typically highly structured and represented in a structured data model such as the relational model. In contrast, the new applications

discussed in this section may require infrequent updates and the queries may be more complex including aggregation and intricate pattern matching queries. In addition, the data may be less structured or completely unstructured. We describe some of the most prevalent of these applications and the underlying DBMS support technologies.

Data Warehouses Data warehouses provide integrated access to historical data collected from legacy data sources [42]. In a typical business, numerous on-line software systems manage and collect data as part of the daily operation of the company. These systems may be transaction processing systems that use a traditional DBMS or they may be specialized applications that squirrel away data in files. The data used by these different applications hold valuable information about past business decisions and outcomes that can be used to improve future decisions. To accomplish this, warehouses integrate the data under a unified schema (structure) and provide access mechanisms that enable efficient use by analysis and decision support packages.

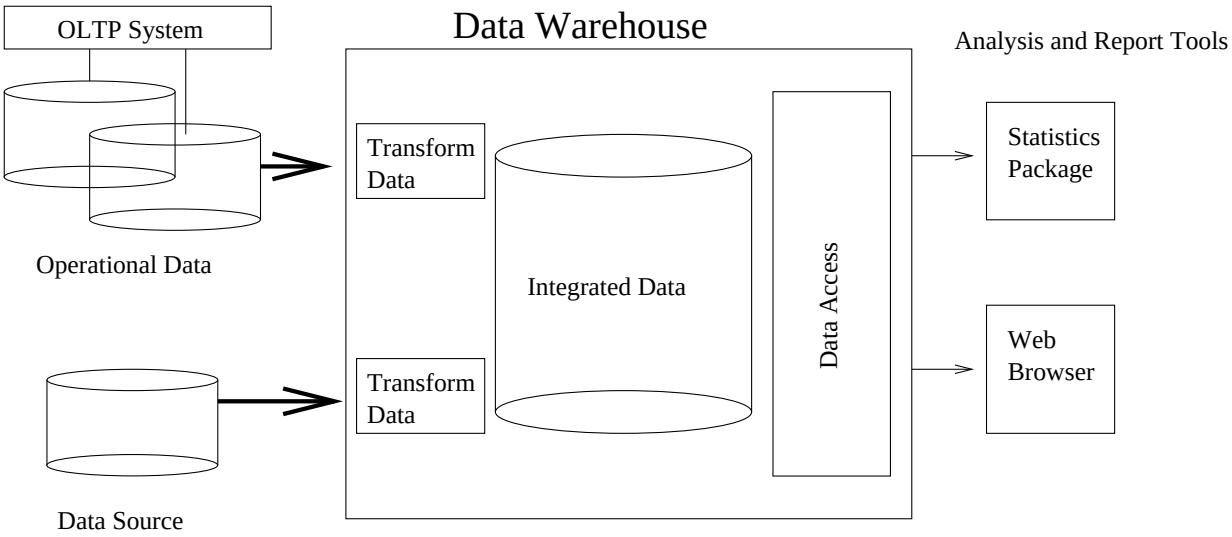

Figure 1: Data Warehouse Architecture.

On-line Analytic Processing (OLAP) OLAP refers to the statistical analysis of data in support of decision making tasks. In OLAP, the focus of data management shifts from one of ensuring consistency and durability of data to one of providing flexible, convenient access to data. As a result, many of the principles that guided the development of data management solutions for OLTP (such as the need to minimize data replication and to normalize data) do not apply to OLAP.

Digital Libraries A digital library is an electronic version of a classical library in which the information resources (for example, books, art work, films) and the indexing information used to locate resources are stored digitally [10]. By its nature, a digital library must be able to store and manage a highly heterogeneous collection of data ranging from unstructured data (for example, images or videos) to semi-structured data (hypertext documents) to structured data (descriptive metadata). Digital libraries use techniques from both information retrieval and structured databases, and extend these with new browsing and searching techniques.

Statistical and Scientific Database Management Systems (SSDBMS) Statistical DBMS are designed to manage socio-economic datasets (for example, census data or economic forecasting data) [36]. Scientific DBMS manage complex collections of data used in and gather from experiments and other scientific endeavors. As in OLAP, SSDBMS must support sophisticated browsing, summarization and analysis functions. In addition, this support must be provided over a diverse collection of complex data, including not only numeric and text data, but also data with complex types. These types may represent such objects as molecular structures, terrain maps, or architectural plans.

World-Wide-Web and Databases The use of DBMS to store World-Wide-Web (Web) content has proven to be an effective means of creating dynamic, scalable Web servers. Using interfaces such as the common gateway interface (CGI), Web application programs can access DBMS to retrieve static Web pages or to dynamically create pages based on query results. While DBMS may be used to store

web content, the Web also permits the electronic publishing of existing (or legacy) databases. Users of published databases, unlike traditional DBMS users, are typically unfamiliar with the data and structure of the database. As a result, users may be unable to effectively formulate structured queries and require new solutions for browsing and effectively locating data in large, complex datasets [23].

To meet the data management needs of these emerging applications, new support technology has be incorporated into DBMS. We examine this new technology including extensions to data models, query languages, indexing methods, query processing engines and query optimizers.

Data Model

Traditional DBMS use structured data models such as the relational model, hierarchical models or object-oriented models. Structured data models assume data can be grouped into collections of tables or classes each having a well defined structure (or schema). To accommodate the needs of new applications, data models have been extended in three primary directions: direct support for abstract data types; addition of conceptual structures to help in the summarization and browsing of large, complex data collections; and support for unstructured and semi-structured data. We examine each of these extensions in turn.

Abstract Data Types Traditional DBMS support a fixed set of simple data types (for example, integers and dates). Extensible DBMS can be extended dynamically with user-defined types and functions. These types can be used to model complex objects, for example molecular structures, along with the behavior of these objects. Most commercial DBMS (including Informix, DB2 and Oracle) now provide such extensibility. To fully support these new data types, a DBMS must provide data management support including new indexing and query processing techniques.

Multi-dimensional Models In data warehousing, OLAP, and statistical applications, data is often conceptually modeled as having multiple dimensions. For example, census data can be viewed as having

the dimensions profession, age, years-at-current-address, *etc.* Product design data can be viewed as having dimensions designer, product type, date-of-production, *etc.* In these examples, the tables containing census or product data are called fact tables. Such a multi-dimensional view of data facilitates the direct modeling of potentially complex relationships among dimensions. For example, the date-of-production dimension may be refined into subdimensions day, month, and year. There is a functional relationship from date-of-production to each of its subdimensions. A multi-dimensional model also facilitates the expression of aggregation or summarization of data along different dimensions (or subdimensions) of interest to a user. For example, a user may retrieve the number of people over age 35 in each state who have technology related professions. The explicit modeling of dimensions provides a convenient formalism on which language operators for aggregation and summarization can be built.

Unstructured and Semi-Structured Models Traditional Information Retrieval (IR) systems use unstructured data models to represent data. Data are stored in *documents* of arbitrary type and structure. Hence, documents may be images, video sequences or full text data stored in any format. Each document is modeled as a bag of words (which may be a subset of the words in a document or a set of words describing an image or video). No structure is associated with these words so a document may contain the word `Washington`, but the model does not include information on whether `Washington` is the author, the subject, or the location of the document. Unstructured models are appropriate for data that truly has no inherent structure. However, they fail to provide sufficient functionality when used to model data (such a Web pages) that have some structure. Consider an XML document which may have tags indicating the author, creation date, and title of the document along with large portions of unstructured data (such as the body of the document). Using an unstructured data model, a query could not be posed to retrieve Web documents written by `Washington`. Using a structured data model to represent this information is equally unsatisfying since Web documents rarely share the same structure. At best one could define a table containing attributes common to most documents. To provide better

support for such data, semi-structured data models have been developed. These models are often self-describing data models in which a data object is described by both a value and its structure. Hence, each object may have its own unique structure. In addition, these models often permit objects to be associated with other objects, typically using labeled graphs [5].

Query Language

In structured data models, the structure (or schema) is used as the primary vehicle for querying. In structured query languages (such as SQL or OQL), schema components (for example, attribute, relation or class names) are used to specify what data should be retrieved. Hence, the user must know and understand the schema in order to pose queries. In unstructured data models, the query model is based on keyword matching. A set or boolean combination of user specified keywords are matched against the words representing the stored documents. To support efficient querying, indexes such as inverted indexes are used to quickly map keyword(s) to documents which they annotate. Sophisticated techniques are used to ensure that all relevant documents are retrieved and no irrelevant documents. These techniques include linguistic techniques for detecting synonyms among keywords.

Query languages for semi-structured data models permit the specification of structured queries over data objects with known structure. However, given that each data object may have its own structure, understanding the structure of an entire database may not be feasible. In a data warehouse, where the structure of the data may be extremely complex, users may need to pose queries without knowing the full structure. As a result, semi-structured query languages permit the specification of pattern matching style queries (e.g., "Find all building plans designed by Maria that include a heating system with more than 100 subcomponents."). Such queries permit the browsing and location of data in unknown or partially known structures [1].

Pattern matching queries are also useful in querying heterogeneous structures. Multi-database languages provide additional data restructuring and merging operations to facilitate data integration.

Meta-data, that is descriptive information about database schemas, can be extremely valuable in enabling the integration of heterogeneous data sources. Higher order languages that permit the querying of schemas along with the data have been used successfully in heterogeneous DBMS.

Data warehouses provide powerful aggregation and summarization facilities to permit the extraction of relevant information. The aggregation functions typically include the basic functions provided in SQL (and object-based variants of SQL) for computing counts, averages, sums, maximums and minimums along with more sophisticated statistical functions over numeric data. Some DBMS permit the user to define new aggregation functions. The summarization techniques extend the simple horizontal partitioning permitted by the SQL *group by* operator. The group by operator partitions the tuples of a table (or instances of a class) into groups based on the values of a set of attributes. Aggregation functions are then applied to each group to compute a summary for each set of attribute values. Extensions permit the partitioning to be based on the values of any function applied to the table attributes. The *cube* operator is used to compute cross tabulations on a table [18]. In contrast to the group by operation, a cross tabulation includes sub-totals (or sub-aggregates) for every subset of attributes.

Browsing techniques provide a convenient way of introducing the data and schema to new users. Users can navigate through the data, effectively locating data of interest. Browsing techniques can be broadly grouped into two strategies. The first uses concept classifications from library science to logically organize the database. Documents are associated with concepts in the classification, such as agriculture or welding. Concepts are related to each other based on their semantic relationships. Users can browse the concept classification, which is often presented using hypertext, to locate documents. The second type of browsing technique uses OLAP style summarizations of the database to permit users to locate data of interest [23]. These systems group together subsets of the database and present aggregates of the data items in each group. Hierarchical abstractions, or dimensions, over the data are used to form the aggregates. A user may *drill-down* into a set of data by successively restricting one or more of the dimensions, while the system presents aggregates of the underlying data at each step.

Indexing

New data intensive applications require much more complex forms of querying. This complexity can take on many forms. In semi-structured data, path queries, which retrieve a subset of objects directly or indirectly associated with a give object, are common. In OLAP and SSDBMS, aggregate queries and multi-dimensional queries are common. In many of these application, queries may join multiple tables (for example, a fact table may be joined with many dimension tables). These query characteristics require the development of new indexing mechanisms for enhancing query performance [27, 40].

To support complex queries over multiple tables, multi-table indices have been developed including join indices and star indices. These indices materialize (i.e., cache) common joins enabling complex queries to be performed efficiently. Traditional indices have been generalized to enable the indexing of new user-defined types. To support queries with independent selection conditions, bitmap indices may be used. A bitmap index is a modification of a traditional index (for example, a B^+-tree or hash index) where for each index value (or key), the index stores a bitmap representing which tuples contain the given value. Bitmap indices have also been shown to be useful in enhancing the execution of some aggregate queries. Other specialized access structures are tailored to materialize specific, commonly used queries (for example, projection indices). The update characteristics of these structures may be unacceptable for OLTP applications. However, for read-only or read-mostly applications, the improved query speed may offset any additional update cost. Data may also be replicated and stored under materialized views. Many data management products make extensive use of materialized aggregate views, including materialization of the data cube, to permit fast computation of aggregate summaries.

Query Processing

The query language extensions and new indexing structures outlined above introduce a variety of new challenges for query processing. The proliferation of new physical structures for accessing data has required the development of new techniques for determining when an index structure or materialized

view can be used (correctly) in answering a query [19]. Similar techniques have also been applied to heterogeneous DBMS to enable query processing over heterogeneous views of data [6, 20, 33, 38]. New efficient algorithms for computing the data cube and other aggregate queries have been developed and incorporated into commercial query engines.

Query Optimization

To complement the new query processing strategies, new techniques for query optimization have been required. Given the new language operators and the new access methods available, the task of deciding which combination of operators and which indices or view to use in executing a query has become significantly more difficult. Query optimization is already a complex task in conventional relational systems. The challenge for new applications is to introduce new operators and access structures in a way that does adversely effect the performance or quality of the query optimizer.

Recent research has addressed some of the issues involved in optimizing aggregate queries and queries with expensive (possibly user-defined) functions. Magic sets, and their cost-based extensions, have proven valuable in optimizing complex relational queries, including queries over views [35]. Algebraic and cost-based optimization of queries over heterogeneous DBMS has also been addressed, though much work remains to be done (see [39] for a summary). Work on optimizing queries over semi-structured data has just begun.

Multimedia Database Management Systems

Recent advances in computing, communication, and storage technologies have enabled the proliferation of *multimedia data*, such as images, graphics, audio, video, and animation, in a number of diverse application domains. Examples of such domains include digital libraries, satellite image archival and processing, training and education, entertainment, and medical databases containing X-rays and MRIs. Currently, the bulk of multimedia data resides in either conventional or multimedia storage servers

offering special-purpose, application-specific functionality. This situation, however, raises a number of problems including redundancy, inconsistency, concurrent access anomalies, as well as integrity, atomicity, and security problems. The continuously expanding diversity of multimedia applications and volume of multimedia data further exacerbate the problem. Incorporating database technology in multimedia application development can offer several benefits, including declarative query facilities, transparency from physical aspects of storage, associative access through indexing, data consistency through well-defined access methods, multi-user access through concurrency control, and reliability through recovery mechanisms. This understanding has given rise to a significant amount of recent interest in *multimedia database management systems* [3, 12, 26, 31, 37].

Providing database functionality for multimedia data types presents a host of new challenges not addressed by conventional DBMS. These challenges stem from the fact that multimedia data types differ from traditional alphanumeric data in their characteristics, and hence require different techniques for their organization and management. A first distinguishing characteristic of multimedia data is their *volume* – a JPEG-compressed color image can require several Megabytes of storage space and a 100 minute video compressed using the MPEG-I standard requires about 1.25 Gigabytes of storage space. Conventional DBMS and file systems provide only very limited support for such large objects, typically in the form of special data types, such as *long fields* and *BLOBs* (Binary Large OBjects), with very poor semantics. Reducing multimedia data to single, large, uninterpreted data values is clearly inadequate for supporting the rich semantic content of multimedia data types and places the whole burden of data processing within the application. A second, and perhaps most important, characteristic of multimedia data types is that, in contrast to alphanumeric data, they are typically characterized by a *spatial extent* (e.g., images and graphics), a *temporal extent* (e.g., audio and speech), or *both* (e.g., video). As a consequence, multimedia data have much richer semantics than conventional symbolic data and any meaningful interpretation of a multimedia object is typically based on its relationship to a system of spatial coordinates and/or a constantly progressing time scale. Furthermore, time-dependent

multimedia data (also known as *Continuous Media* (CM) data) like audio and video have specific *timeliness* constraints associated with them. For example, a video clip consists of a *stream* of video frames which must be delivered to viewers at a certain rate (typically 30 frames/sec). For MPEG-I compressed video, this translates to a data rate of approximately 1.5 Megabits per second (Mbps). The underlying storage manager needs to ensure that the storage and retrieval of CM data proceeds at their pre-specified real-time rates [29]. Integrated support for the spatio/temporal nature and semantics of multimedia data requires non-trivial extensions to various basic building blocks and functional units of a DBMS.

Data Model

Complex multimedia objects require sophisticated modeling mechanisms with rich semantic capabilities. An important requirement for these conceptual tools is the ability to model the complex spatio/temporal structure of a multimedia object through well-defined abstractions. For spatial non-continuous data, like images, the modeling problem is probably simpler since the semantics of objects and operations are clearly defined and their properties can be derived from geometry. CM data, on the other hand, presents the much more difficult problem of modeling *time* with conceptual mechanisms that can capture: (1) *Intra-Media Continuity*, that is, the real-time delivery requirements of a CM stream; (2) *Inter-Media Synchronization*, that is, the precedence and real-time synchronization constraints among the component CM streams of a complex multimedia presentation (e.g., audio and video "lip-synch"ing); and, (3) *User Interaction*, that is, the ability of a user to interact with the presentation through standard VCR-type functions (e.g., fast-forward or rewind) that can change the presentation speed or randomly access specific points in a presentation.

Most efforts for managing multimedia data have been based on flexible object-oriented or extended relational models that allow for the modeling of complex structured multimedia objects, the definition of abstract media types and operations on media data units. However, despite their ability to model

complex structures such data models lack the temporal modeling capabilities required by CM data – the problems of stream-oriented, real-time access and synchronization still remain. A number of conceptual models have been developed for capturing the temporal aspects of multimedia data. They can be roughly classified into three categories, namely: *graph-based models* (e.g., object composition petri nets [21] and presentation graphs [26]), *language-based models* (e.g., HyTime [25] and MHEG [34]), and *temporal abstraction models* (e.g., temporal intervals and relations [2] and timed streams [17]). Nevertheless, the efficient mapping of such conceptual constructs to the physical level of a full-fledged multimedia DBMS still remains an issue of concern.

Query Language

Declarative query languages are an important part of DBMS and have played an important role in their success. A powerful declarative querying facility allows associative (i.e., content-based) access to the underlying data and helps to maintain the desired independence between the DBMS and the application. Conventional DBMS query languages are typically based on the assumption of highly symbolic alphanumeric representations and, thus, cannot accommodate the much richer spatio/temporal semantics of multimedia data. More specifically, query languages for complex multimedia objects need to address the following issues:

1. *Similarity Queries.* Conventional declarative content-based querying is based on *exact-matching* between well-defined sets of symbols using simple equality or comparison operators. An example of such a query is: "Select all employees with salary $>$ 45K". For any employee in the database, the search condition will evaluate to either `TRUE` or `FALSE`, based on a well-understood numerical comparison. Such exact-matches are rarely of interest for multimedia data types such as images or video. Users are usually interested in discovering multimedia objects that are perceptually *similar* (to each other or to some query object), where the notion of similarity typically depends on the data type and the requirements of the application. Answers to such similarity queries will

be *ranked* based on grades of similarity obtained using an appropriate similarity function and users will usually be interested in obtaining the `TOP-k` results, i.e., the objects with the `k` highest grades [8, 13].

2. *Spatio/Temporal Queries.* A complete declarative query facility for multimedia DBMS should allow users queries not only on the content, but also the *structure*, i.e., the spatio/temporal characteristics of multimedia objects. Examples of such queries include spatial searches (e.g., "Find all the roads passing through Murray Hill"), temporal searches (e.g., "Find all scenes where President Clinton is shaking hands *after* stepping off an airplane"), and simple spatial or temporal computations (e.g., "Find the area of this object"). Of course, supporting spatio/temporal queries is intimately related to the facilities offered by the underlying data model for modeling the complex spatial and temporal structure of multimedia objects.

3. *Quality of Service (QoS) Specifications.* Multimedia objects can often be accessed at multiple levels of resolution or *Quality of Service (QoS)* that correspond to different service requirements on the underlying DBMS resources. Important QoS parameters include the average delay (experienced by the user), the actual presentation rate and image resolution, and the allowable deviations for temporal synchronization constraints. Some application environments can be flexible about certain QoS parameters (e.g., audio quality or image resolution). Since such flexibilities can be directly translated to flexible resource requirements, effective QoS specifications play a very important role in effective query processing and optimization in a multimedia DBMS [7].

Indexing

Similarity-based queries are the prominent form of associative data access in a multimedia DBMS. Efficient execution of such queries requires the development of appropriate indexing mechanisms for retrieval by similarity. The standard technique for this purpose is to map both the query and each

multimedia object into some multi-dimensional *feature space* such that two perceptually similar objects are guaranteed not to be far apart in this space [12]. Typical features of multimedia objects include color, texture (e.g., contrast, coarseness), shape, text (i.e., a set of keywords or annotations), and motion. There can also be some features specific to particular application domains. Features are extracted either manually or using automated (usually domain-specific) methods, and stored as a collection of feature vectors in the database. For example, the QBIC (Query By Image Content) system developed at IBM Almaden supports queries based on example images, user sketches and drawings, color, texture, shape, and keywords [13]. Color-based querying is implemented by storing a *color histogram* for each image in the database and comparing the color histogram of the query image to those in the database.

Mapping objects and queries onto feature vectors enables the use of appropriate multi-dimensional indexing mechanisms such as grid files and R-trees, with the query region appropriately expanded around the query point. Given that notions of similarity are in general diverse and application dependent, it is important to select appropriate *distance measures* in the multi-dimensional feature space that closely match the perception of similarity. For example, the distance between color histograms in QBIC is defined as a quadratic form function (a generalization of Euclidean distance) that takes into account the "cross-talk" between two similar colors (e.g., orange and red) [13]. One issue that needs to be addressed is that conventional multi-dimensional indexing methods like grid files or R-trees suffer from the infamous "dimensionality curse", meaning that they result in access times that are exponential in the number of dimensions or they degrade to a linear search as the dimensionality increases. This is a serious problem for multimedia data indexing, since the number of dimensions (i.e., features) can in some cases exceed one hundred [41]. One approach for dealing with high dimensionality is to map high dimensional feature vectors to a lower number of dimensions using a distance-preserving transformation [12]. Another approach is to design new, scalable indexing structures or to improve existing ones to scale to higher dimensions [41].

Query Processing

Multimedia data types introduce a host of new challenges for the query processing component of a DBMS. A central issue is that the real-time access characteristics and the large volumes of CM data mandate the design of *effective resource management strategies* for multimedia query processing. Such strategies should: (1) provide guaranteed service levels for the storage and retrieval of CM data; (2) provide support for the temporal synchronization constraints defined between the CM components of complex multimedia objects; (3) provide support for user interaction (i.e., VCR-type functions); (4) allow for the retrieval of non-continuous data concurrently with CM data; and, (5) maximize system throughput and reduce system response times. A number of these issues have been addressed in the context of CM storage servers, e.g., the *Fellini* multimedia storage server developed at Bell Labs [22].

Given the limited amount of DBMS resources (e.g., memory, disk bandwidth, disk storage), providing service guarantees for CM data mandates a novel *admission control* component that decides whether to execute or postpone user queries. By initiating the execution of a query, the DBMS commits to satisfy the resource requirements (e.g., memory, disk bandwidth) of the CM streams involved throughout their duration. The service guarantees provided by the admission control policy can be either deterministic (i.e., based on worst-case assumptions) [29] or stochastic (i.e., based on statistical models of system behavior) [24]. Prior research has proposed novel data layout strategies, disk scheduling algorithms, and buffer management policies that take advantage of the highly sequential, stream-oriented access patterns to CM data in order to improve system throughput [4, 32]. A method proposed for handling conventional (non-continuous) data requests and user interaction is to reserve a portion of the system's resources specifically for that purpose [22]. Given that typical CM requests tend to execute for long periods of time, reserving resources in advance is important to ensure that both conventional requests and VCR-type functions observe reasonable response times. Other schemes for implementing VCR-type functionality have also been proposed in the literature, for example, storing a fast-forward/rewind version of a CM stream [30]. Finally, efficient resource scheduling algorithms for complex multimedia presentations with

user-defined synchronization between various CM streams have recently been proposed [14].

Another crucial problem is the design of efficient query processing strategies for handling similarity queries. Given that users are interested only in the `TOP-k` objects, new query execution methods are needed to produce the best `k` results efficiently (i.e., without materializing every intermediate result object that matches at all). Efficient strategies for processing `TOP-k` queries have been investigated in the context of the Garlic project at IBM Almaden [6, 11]. Other important query processing issues for a multimedia DBMS include effective handling of tertiary storage and hierarchical storage structures (given the voluminous nature of multimedia data) [16], techniques for *sharing* CM streams among users to reduce resource demands [9, 15], and fault-tolerant operation [28].

Query Optimization

The declarative query language interface offered by the majority of conventional database systems has definitely been a major factor in their commercial success. This declarative querying paradigm was made possible due to the development of sophisticated query optimization techniques. Likewise, supporting high-level queries over multimedia databases mandates the development of appropriate optimization techniques. Multimedia query optimization is still a very open research area with most important problems still waiting to be formulated and adequately solved.

The major issue here is that for a multimedia DBMS the querying model and, therefore, the resulting optimization questions differ in many ways from conventional DBMS querying. Recent research has addressed some of the issues involved in optimizing similarity-based selection queries with ranked result sets over multimedia repositories [8]. The main emphasis in this work was to explore optimization strategies designed for graded results and `TOP-k` semantics. Two additional issues that arise in the optimization of multimedia queries are Intra/Inter-Media Synchronization and QoS [7]. Ignoring synchronization constraints during optimization can lead to excessive buffer requirements and underutilization of resources at run-time or unacceptable flaws in the presentation (e.g., glitches in the video,

out-of-sync audio). QoS requirements are significant for optimization since they impact the space of execution alternatives as well as the metric of optimization. For example, a query generated by a fraud detection application needs to be evaluated speedily with quality of video being of secondary importance. Thus, the optimizer should obviously consider the option of returning a low-quality (e.g., compressed) version of the video if this results in lower response time. To the best of our knowledge, these issues have yet to be addressed by the database or multimedia research community.

References

[1] S. Abiteboul. "Querying Semi-Structured Data". In *Proceedings of the Sixth International Conference on Database Theory (ICDT'97)*, pages 1–18, 1997.

[2] J.F. Allen. "Maintaining Knowledge About Temporal Intervals". *Communications of the ACM*, 26(11):832–843, November 1983.

[3] P.M.G. Apers, H.M. Blanken, and M.A.W. Houtsma, editors. *"Multimedia Databases in Perspective"*. Springer-Verlag, 1997.

[4] Steven Berson, Shahram Ghandeharizadeh, Richard Muntz, and Xiangyu Ju. "Staggered Striping in Multimedia Information Systems". In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data*, pages 79–90, Minneapolis, Minnesota, May 1994.

[5] P. Buneman. "semi-structured data". In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, 1997.

[6] Michael J. Carey, Laura M. Haas, Manish Schwarz, Peter M. Arya, William F. Cody, Ronald Fagin, Myron Flickner, Allen W. Luniewski, Wayne Niblack, Dragutin Petkovic, John Thomas, John H. Williams, and Edward L. Wimmers. "Towards Heterogeneous Multimedia Information Systems: The Garlic Approach". In *Proceedings of the Fifth International Workshop on Research Issues in Data Engineering - Distributed Object Management (RIDE-DOM'95)*, pages 124–131, Taipei, Taiwan, March 1995.

[7] Surajit Chaudhuri. "On Optimization of Multimedia Queries". In *Proceedings of the ACM Multimedia '94 Conference Workshop on Multimedia Database Management Systems*, San Francisco, California, 1994.

[8] Surajit Chaudhuri and Luis Gravano. "Optimizing Queries over Multimedia Repositories". In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pages 91–102, Montreal, Canada, June 1996.

[9] Asit Dan, Dinkar Sitaram, and Perwez Shahabuddin. "Scheduling Policies for an On-Demand Video Server with Batching". In *Proceedings of ACM Multimedia '94*, pages 15–23, San Francisco, California, October 1994.

[10] *Communications of the ACM: Special Issue on Digital Libraries*, volume 38. ACM Press, April 1995.

[11] Ronald Fagin. "Fuzzy Queries in Multimedia Database Systems". In *Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, Seattle, Washington, June 1998. (Invited Paper).

[12] Christos Faloutsos. *"Searching Multimedia Databases by Content"*. Kluwer Academic Publishers, 1996.

[13] Christos Faloutsos, Ron Barber, Myron Flickner, J. Hafner, Wayne Niblack, Dragutin Petkovic, and William Equitz. "Efficient and Effective Querying by Image Content". *Journal of Intelligent Information Systems*, 3(3/4):231–262, July 1994.

[14] Minos N. Garofalakis, Yannis E. Ioannidis, and Banu Özden. "Resource Scheduling for Composite Multimedia Objects". In *Proceedings of the 24th International Conference on Very Large Data Bases*, New York City, U.S.A., August 1998.

[15] Minos N. Garofalakis, Banu Özden, and Avi Silberschatz. "Resource Scheduling in Enhanced Pay-Per-View Continuous Media Databases". In *Proceedings of the 23rd International Conference on Very Large Data Bases*, pages 516–525, Athens, Greece, August 1997.

[16] Shahram Ghandeharizadeh and Cyrus Shahabi. "On Multimedia Repositories, Personal Computers, and Hierarchical Storage Systems". In *Proceedings of ACM Multimedia '94*, pages 407–416, San Francisco, California, October 1994.

[17] Simon Gibbs, Christian Breiteneder, and Dennis Tsichritzis. "Data Modeling of Time-Based Media". In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data*, pages 91–102, Minneapolis, Minnesota, May 1994.

[18] J. Gray, A. Bosworth, A. Layman, and H. Pirahesh. "Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals". In *Proceedings of the Twelfth International Conference on Data Engineering*, pages 152–159, 1996.

[19] A. Y. Levy, A. O. Mendelzon, Y. Sagiv, and D. Srivastava. "Answering Queries Using Views". In *Proceedings of the Fourteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, San Jose, CA, May 1995.

[20] A. Y. Levy, D. Srivastava, and T. Kirk. "Data Model and Query Evaluation in Global Information Systems". *Journal of Intelligent Information Systems: Special Issue on Networked Information Discovery and Retrieval*, 5:121–143, 1995.

[21] Thomas D.C. Little and Arif Ghafoor. "Sunchronization and Storage Models for Multimedia Objects". *IEEE Journal on Selected Areas in Communications*, 8(3):413–427, April 1990.

[22] Cliff Martin, P.S. Narayanan, Banu Özden, Rajeev Rastogi, and Avi Silberschatz. "The $\mathcal{F}_{ellini}$ Multimedia Storage Server". In *"Multimedia Information Storage and Management", S.M. Chung (Ed.)*. Kluwer Academic Publishers, 1996.

[23] R. J. Miller, O. G. Tsatalos, and J. H. Williams. "DataWeb: Customizable Database Publishing for the Web". *IEEE Multimedia*, 4(4):14–21, Oct-Dec 1997.

[24] Guido Nerjes, Peter Muth, and Gerhard Weikum. "Stochastic Service Guarantees for Continuous Media Data on Multi-Zone Disks". In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 154–160, Tucson, Arizona, May 1997.

[25] Steven R. Newcomb, Neill A. Kipp, and Victoria T. Newcomb. "The `HyTime` Hypermedia/Time-based Document Structuring Language". *Communications of the ACM*, 34(11), November 1991.

[26] Kingsley C. Nwosu, Bhavani Thuraisingham, and P. Bruce Berra, editors. *"Multimedia Database Systems: Design and Implementation Strategies"*. Kluwer Academic Publishers, 1996.

[27] P. O'Neill and D. Quass. "Improved Query Performance with Variant Indexes". In *Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data*, Tuscon, Arizona, June 1997.

[28] Banu Özden, Rajeev Rastogi, Prasant Shenoy, and Avi Silberschatz. "Fault-tolerant Architectures for Continuous Media Servers". In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pages 79–90, Montreal, Canada, June 1996.

[29] Banu Özden, Rajeev Rastogi, and Avi Silberschatz. "A Framework for the Storage and Retrieval of Continuous Media Data". In *Proceedings of the 1995 International Conference on Multimedia Computing and Systems*, pages 2–13, Washington, D.C., May 1995.

[30] Banu Özden, Rajeev Rastogi, and Avi Silberschatz. "On the Design of a Low-Cost Video-on-Demand Storage System". *ACM Multimedia Systems*, 4(1):40–54, 1996.

[31] Banu Özden, Rajeev Rastogi, and Avi Silberschatz. "Multimedia Support for Databases". In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, Tucson, Arizona, May 1997.

[32] Banu Özden, Rajeev Rastogi, Avi Silberschatz, and Cliff Martin. "Demand Paging for Movie-on-Demand Servers". In *Proceedings of the 1995 International Conference on Multimedia Computing and Systems*, Washington, D.C., May 1995.

[33] Y. Papakonstantinou, H. Garcia-Molina, and J. Widom. "Object Exchange Across Heterogeneous Information Sources". In *Proceedings of the Eleventh International Conference on Data Engineering*, Taipei, Taiwan, March 1995.

[34] R. Price. "MHEG: An Introduction to the Future International Standard for Hypermedia Object Interchange". In *Proceedings of ACM Multimedia '93*, pages 121–128, Anaheim, California, August 1993.

[35] P. Seshadri, J. Hellerstein, H. Pirahesh, T.Y. C. Leung, R. Ramakrishnan, D. Srivastava, P. J. Stuckey, and S. Sudarshan. "Cost-Based Optimization for Magic: Algebra and Implementation". In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pages 435–446, Montreal, Canada, June 1996.

[36] A. Shoshani. "OLAP and Statistical Databases: Similarities and Differences". In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 185–196, Tuscon, Arizona, May 1997.

[37] V.S. Subrahmanian and Sushil Jajodia, editors. *"Multimedia Database Systems: Issues and Research Directions"*. Springer-Verlag, 1996.

[38] D. Suciu. "Query Decomposition and View Maintenance for Query Languages for Unstructured Data". In *Proceedings of the 22nd International Conference on Very Large Data Bases*, pages 227–238, 1996.

[39] A. Tomasic, L. Raschid, and P. Valduriez. "Scaling Heterogeneous Databases and the Design of DISCO". In *Proceedings of the International Conference on Distributed Computer Systems*, 1996.

[40] O. Tsatalos, M. Solomon, and Y. Ioannidis. "The GMAP: A Versatile Tool for Physical Data Independence". *The VLDB Journal*, 5(2), April 1996.

[41] David A. White and Ramesh Jain. "Similarity Indexing with the SS-Tree". In *Proceedings of the Twelfth International Conference on Data Engineering*, pages 516–523, New Orleans, Louisiana, February 1996.

[42] J. Widom. "Research Problems in Data Warehousing". In *Proceedings of the International Conference on Information and Knowledge Management*, November 1995.