GUEST EDITORS' INTRODUCTION TO PART I OF THE SPECIAL SECTION

Data and Knowledge Management in Multimedia Systems

P. Bruce Berra, *Fellow, IEEE*, and Arif Ghafoor, *Senior Member, IEEE*

1 Introduction

Emerging multimedia technologies will allow users to store, retrieve, share, and manipulate complex information in order to build exciting multimedia applications. This will require substantial changes in our approach to the design of conventional database and storage systems.

Multimedia database management systems are necessary for a wide variety of commercial, governmental, and personal applications. These include telemedicine, distance learning, video surveillance, law enforcement, traffic management, GIS, video-on-demand, digital libraries, etc. Most of these applications require the use of large multimedia data repositories.

The stored nature of the information poses a number of challenges in the management of multimedia information, including data and knowledge representation, indexing and retrieval, integration, intelligent searching techniques, information browsing, and query processing.

Several technical challenges are faced by the research community that need to be tackled in order to provide viable solutions for the development of multimedia database systems. For example, digital libraries of the future will have massive amounts of documents. Intelligent techniques are needed to allow efficient and automatic generation of indices and cataloging of multimedia documents based on user-specified object and content descriptions. Emerging telemedicine technology is aimed at automating the storage and retrieval of patient records consisting of different types of multimedia data, including medical records, radiological data, physician's briefs, etc.

In a distributed collaborative manufacturing enterprise, sharing and retrieval of pertinent multimedia data related to different aspects of industrial manufacturing and product maintenance is critical to the productivity and growth of the enterprises. For all such applications, multimedia database systems must exhibit high degrees of robustness, efficiency, and reliability for deployment in real-world environments.

Key challenges in the development of multimedia database systems include the following:

- content analysis of images and video data to extract features for recognition of salient objects and faces;
- knowledge-based representation of multimedia data;
- semantic modeling of events of interest and management of spatio-temporal metadata;
- facilities to retrieve multimedia data containing contents of interest to users, including query interface to enable users to formulate semantic queries in an intuitive manner; and
- design and development of high-performance multimedia database servers for networked environments.

Automation of content-based retrieval in images and video databases requires detailed identification of objects and events in the data. Traditionally, research in image and video database systems has focused on the development of robust image processing and object recognition techniques.

Image contents are generally classifiable at several levels of granularity, ranging from fine grained, which focus on minute details, to coarse grained, which deal with larger perspectives of image semantics. Management and retrieval of images at different levels of detail entails several challenges, including design of robust image processing techniques for feature extraction, development of efficient data models for content-based indexing and retrieval, knowledge-based representations of objects, and domain-specific semantic modeling of events of interest. General approaches to address these issues are mainly focused on developing multilevel abstraction mechanisms. Advanced multifeature models are necessary to overcome the hurdles in robust and accurate human face and object recognition.

The handling of imprecision in the representation of knowledge and data related to images is another open challenge. Imprecision in query formulation compounds the problem of accurate retrieval of images. Intelligent query formulation mechanisms are needed, which allow

- *P.B. Berra is with the Department of Computer Science and Engineering, Wright State University, Dayton, OH 45435. E-mail: bberra@cs.wright.edu.*
- *A. Ghafoor is with the School of Electrical and Computer Engineering, Purdue University, 1285 Electrical Engineering Building, West Lafayette, IN 47907. E-mail: ghafoor@ecn.purdue.edu.*

For information on obtaining reprints of this article, please send e-mail to: tkde@computer.org, and reference IEEECS Log Number 108124.

the posing of fuzzy queries and ensure a high degree of accuracy in image retrieval.

In the context of video databases, the problem of indexing and retrieval is even more difficult. Semantics and events in video data can be expressed by describing the interplay among physical objects over time, along with spatial relationships between these objects. An important open problem is the unambiguous modeling of such interplay. Several knowledge-based techniques, such as spatio-temporal logic, semantic network models, etc., can be utilized for conceptualizing video events. Event identification and indexing of spatio-temporal semantics requires processing of a sequence of video frames to extract motion information. In essence, powerful mechanisms are needed for modeling object motions and their relationships in time and space. These mechanisms can subsequently lead to suitable indexing structures used in query processing.

Another critical issue is how to deal with the heterogeneity that may exist among semantics of multimedia data due to difference in the preconceived interpretation or intended use of the information given in an image or in a sequence of video frames by different users. Semantic heterogeneity is a difficult problem even for conventional databases, with little or no consensus on the way to tackle it in practice.

The emergence of multimedia applications focusing on large archives of electronic documents is generating interesting challenges in the design of high-performance multimedia document servers. With the rapid proliferation of the World Wide Web and continuous advances in broadband networking technology, there is a growing trend toward large-scale multimedia document production and archive systems. The fundamental technical issue in managing such archives is the synchronization of several different media types together in order to present information to the end user in an integrated form. Synchronization of multimedia data requires the enforcement of spatio-temporal constraints among various media. Models are required for specification of such synchronization. These models are generally integrated within the database as part of metaschema. Synchronization problem of multimedia data becomes even more challenging when data is transported over various networks with varying characteristics, which is also true for the Web.

Storage management of multimedia data archives, in particular video data, is an important technological challenge in the development of viable large-scale multimedia servers. Multimedia storage servers form a core technology in providing access to multimedia data. These servers must allow real-time storage and retrieval requirements despite the large storage space and data transfer rates of digital audio, images, and video. The multimedia research community has been grappling with these problems by developing architectures and algorithms for efficient high-performance digital multimedia servers.

Another area of research being pursued by the multimedia research community is the provision of intuitive and flexible user interfaces for high-level query formulation for audio and video data. The operations include querying and retrieval of data based on user-supplied attributes. It also includes facilities to abstract, modify, or reprocess archived multimedia data. Facilities for imprecise matching will allow the trade-off between computational complexity and precision in video information retrieval.

2 The Featured Papers

We received 61 papers, of which 12 were selected for publication in this special section. These include 10 regular and two concise papers. The regular papers have been organized into four broad categories:

1) image retrieval,
2) video data modeling,
3) multimedia synchronization, and
4) multimedia storage management.

Due to page constraints, the current section (Part I) contains six regular papers and both concise papers from the first two categories. Papers from the third and fourth categories will appear in Part II in the May/June 1999 issue of *TKDE*.

The four categories of the selected papers depict the wide variety of directions in which research is moving in the multimedia systems area and provide a broad coverage of the issues and challenges discussed above. This special section focuses on several of the open research issues in these areas. Many of the research results presented in this section are supported by significant experimental evaluations. The results reported by the research community hold great potential for solving the above-mentioned technical challenges. These technologies can be further developed to produce viable solutions for these challenges and build mature systems capable of deployment in the real-world.

In "Knowledge-Based Image Retrieval with Spatial and Temporal Constructs," Chu et al. propose a multilayered knowledge-based semantic image model consisting of four layers: the raw data layer, the feature and content layer, the schema layer, and the knowledge layer. Based on this model, they present a system capable of handling certain types of spatial, evolutionary, and temporal queries. In addition, approximate matching of feature and content, conceptual terms, and temporal logic predicates are also supported. Image retrieval based on features and content and associated user-interface techniques presented in this paper, though focused on medical systems in this case, promise to have wide-ranging applicability in other domains.

Korn et al.'s paper entitled "Fast and Effective Retrieval of Medical Tumor Shapes" investigates the problem of retrieving similar shapes from a large database, with an emphasis on medical tumor shapes. They use a natural similarity function, search space pruning techniques, and spatial access methods for indexing and search. Their

techniques lead to fast, efficient, and correct database operations, including range queries and nearest-neighbor queries. Such algorithms and strategies will prove to be crucial in the development of future multimedia systems with extremely large search spaces.

Ortega et al.'s paper, "Supporting Ranked Boolean Similarity Queries in MARS," concentrates on the retrieval subsystem of an image database system and its support for content-based queries. This paper describes how these techniques can be adapted for ranked retrieval over image databases utilizing the Boolean retrieval model. Results of accompanying experiments demonstrate the effectiveness of their model for image-retrieval applications.

In "Data Resource Selection in Distributed Visual Information Systems" by Chang et al., the challenge of how the information can be summarized to enable intelligent visual queries is discussed. They propose an approach of summarizing database information and then selecting databases based on summarized information. Content-based indexing of images is used to record a summary of the visual content, including similarity distributions of the images. The retrieval technique is based on a ranking algorithm using query similarity and features associated with the an object template. Such abstraction techniques are expected to form an important dimension in content-based retrieval of multimedia data.

Content-based retrieval techniques for video data are essential for developing useful multimedia information database systems. Jiang et al.'s paper entitled "WVTDB—A Semantic Content-Based Video Database System on the World Wide Web" describes the design and implementation of a Web-based video database system. A video data model is presented that allows dynamic and incremental video annotation and indexing, multiuser view sharing, and video data reuse. Users can query, retrieve, and browse video data based on their semantic content descriptions and specified temporal constraints. Several techniques are proposed for efficient network bandwidth utilization for this kind of Web-based system. Enhancements proposed in the paper may be extended for effective video summarization as well as implementing knowledge-based and spatio-temporal queries on multimedia data.

Ahanger and Little's paper "Automatic Composition Techniques for Video Production" focuses on automating the processes involved in video composition. This process can be automated if appropriate domain-specific metadata are associated with video segments and composition techniques are established. The authors present techniques to achieve dynamic, real-time, and cohesive video composition and customization. They also identify metrics for evaluation of their techniques with respect to existing manually produced video-based news. These technologies may be ultimately incorporated in the design of highly personalized or domain-specific video applications.

A system for video content analysis suitable for management of a home video library is described by Dimitrova et al. in their concise paper, "Video Content Management in Consumer Devices." The system presents the user with a visual table of contents which provides an overview of the video content and direct access to particular points in the stored video. The system utilizes a computationally inexpensive, simple, and powerful mechanism for cut detection and key frame filtering. The approach outlined in this paper can lead to the development of viable content-based multimedia consumer devices of the future.

Ankerst et al.'s concise paper "A Multistep Approach for Shape Similarity Search in Image Databases" employs neighborhood influence weights to adapt the similarity distance functions to user-specified requirements. The authors' similarity model is based on quadratic forms for which they present a multistep query processing architecture suitable for image databases. They demonstrate an optimal filter selectivity by using an algorithm to reduce the dimensionality of quadratic form-based similarity query results. Experiments on an image database exhibit good scalability. The techniques show promise in efficient processing of very high resolution images.

We hope that the research contributions in this section provide a stimulant to readers to grapple more effectively with pragmatic problems in the development of efficient and powerful multimedia management systems. Such systems are expected to be the basis of tomorrow's knowledge-based society at all levels.

ACKNOWLEDGMENTS

We take this opportunity to express our gratitude to all the authors and referees for their effort and support.

GUEST EDITORS' BIOGRAPHIES

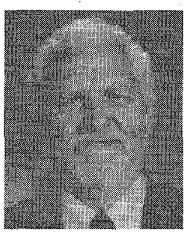

P. Bruce Berra received the BS and MS degrees from the University of Michigan and the PhD degree from Purdue University. He is now director of the Information Technology Research Institute in the College of Engineering and Computer Science and Regents Professor of Information Technology in the Department of Computer Science and Engineering at Wright State University in Dayton, Ohio. He is also an emeritus and research professor of electrical engineering and computer science at Syracuse University. Prior to July 1996, he was director of the New York State Center for Advanced Technology in Computer Applications and Software Engineering (CASE), and a professor of electrical and computer engineering and a faculty member in computer and information science at Syracuse University. He previously served as chair of industrial engineering and operations research at Syracuse University, and has taught at the University of Michigan's Dearborn campus, Boston University, and Purdue University. His industrial experience includes periods of service with Hughes, Bendix, and IBM. He is currently president of PBB Systems, a knowledge and database consulting firm.

Dr. Berra has served on the IEEE Computer Society Board of Governors; was editor-in-chief of the Computer Society Press; was on the IEEE Press Editorial Board and a member of the Computer Society Distinguished Visitors Program; and was general chair and program chair of the IEEE International Conference on Data Engineering. He has served as an editor of *IEEE Transactions on Knowledge and Data Engineering*, *IEEE Transaction on Software Engineering*, *IEEE Transactions on Computers*; and has chaired the IEEE Computer Society Transactions Advisory Committee. He is one of the three co-founders

of *IEEE Transactions on Knowledge and Data Engineering.* He pursues research interests in multimedia information systems, parallel processing for very large data and knowledge bases, and optical database machines. He was recently presented the L.C. Smith College of Engineering and Computer Science Award for Excellence in Scholarship for his contributions to knowledge in his field. In 1994 and 1995, he was a finalist in the Supporter of Entrepreneurship category of the Entrepreneur of the Year Institute Awards. He is a fellow of the IEEE and coguest editor of this section of *TKDE.*

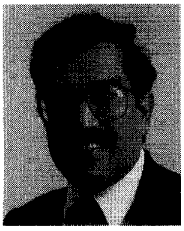

Arif Ghafoor received his BS degree in electrical engineering from the University of Engineering and Technology, Lahore, Pakistan, in 1976; and the MS, MPhil, and PhD degrees from Columbia University in 1977, 1980, and 1984, respectively. In the Spring of 1991, he joined the faculty of the School of Electrical Engineering at Purdue University, where he is now an associate professor. Prior to joining Purdue University, he was on the faculty of Syracuse University from 1984 to 1991. His research interests include design and analysis of parallel and distributed systems, and multimedia information systems. He has published more than 100 technical papers in these areas. Currently, he is directing a research laboratory in distributed multimedia systems at Purdue University. His research in these areas has been funded by DARPA, NSF, NYNEX, AT&T, Intel, IBM, Fuji Electric Corporation, and GE. He has served on the program committees of various IEEE and ACM conferences. Currently, he is on the Editorial Board of *ACM Multimedia Systems*, the *Journal of Multimedia Tools and Applications*, and the *Journal of Parallel and Distributed Databases.* He is also a guest/co-guest editor of special issues of several journals, including *IEEE JSAC*, *ACM Multimedia Systems*, the *Journal of Multimedia Tools and Applications*, and this section of *IEEE Transactions on Knowledge and Data Engineering.* He is a senior member of the IEEE.