2001 NSF CISE/EIA RI and MII PI's Workshop
Yearly Project Report

Purdue University
Multimedia Support Infrastructure Project*

Ahmed K. Elmagarmid, PI
Aref Ghafoor, John T. Korb, Kihong Park, Eugene Spafford

June 15, 2001

1 Introduction

1.1 Purdue University

Purdue University is a coeducational, state-assisted system in Indiana. Founded in 1869 as a land-grant school and named after benefactor John Purdue, the University is one of the nation's leading research institutions with a reputation for excellent and affordable education. Building upon historical strengths in engineering and agriculture, the West Lafayette Campus now offers nearly 6,700 courses in more than 200 specializations in the schools of Agriculture, Consumer and Family Sciences, Education, Engineering, Health Sciences, Liberal Arts, Management, Nursing, Pharmacy and Pharmacal Sciences, Science, Technology, and Veterinary Medicine. The first Department of Computer Sciences in the United States was established at Purdue University in October 1962.

1.2 Multimedia Support Infrastructure Project

Audio and video data form a significant and growing fraction of the total information being created and archived. The diverse nature of this data combined with its magnitude poses difficult problems in organizing, storing, and delivering it in a secure and timely manner. As multimedia information plays an increasingly important role, efforts to solve these problems are becoming critical. Networked multimedia information systems are viewed as catalysts for new research in many areas of computer science and engineering, ranging from basic research to applied technology and education. This view is a result of the fact that no single monolithic service architecture can meet the wide spectrum of characteristics and requirements of various multimedia applications.

Critical issues of managing and translating integrated quality of service requirements across distributed multimedia systems form the unifying theme of the proposed research infrastructure. Strongly coupled to this theme are issues relating to specification, support, and provisions for required levels of quality of service via distributed multimedia servers. When viewed in conjunction with real world constraints and system-wide end-to-end performance requirements, these present formidable research and implementation challenges that encompass all the components of a multimedia support infrastructure (MSI).

The ultimate objective of a comprehensive end-to-end quality-of-service specification relies on several individual component systems such as networks, databases, and security. This research project aims to investigate each of these components with a view to optimizing overall performance. The research infrastructure we are building, is designed to accommodate and advance research in the area of quality of service networking for multimedia databases. Specifically, we are creating a test bed to build and experiment with algorithms and techniques to compress, store, index, retrieve, and deliver multimedia data to applications.

1.3 Project URL

The MSI project web page is available at `<http://www.cs.purdue.edu/msi>`.

* This material is based upon work supported by the National Science Foundation under Grant No. 9972883. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

2 Accomplishments

2.1 Goals, Objectives, and Targeted Activities for the Past Year

We are pleased to report great strides in several areas. We have ordered most of the equipment on schedule for the first year and a half, we have been successful in getting the proper staffing, and we continue to involve faculty and students from historically black colleges and universities in our research activities. We continue to be successful in leveraging the NSF CISE RI funds in several significant and noticeable ways and have been extremely successful in obtaining research funds and grants from numerous sources. Noted among them is our recently awarded grant from the State of Indiana for supporting faculty and graduate students to conduct research and experimental development of a multimedia system for the telemedicine application. These funds are in direct support of this effort. The total funding from the State is over $1.7 million for two years, starting February 1, 2000.

In addition, we have received the first year of cost sharing promised by Telcordia and we are in the process of obtaining the second year of funding. We have also established research collaboration in the area of multimedia document delivery with Siemens Corporation. We have received internal funding from the Center for Education and Research in Information Assurance and Security (CERIAS) to support one post doctoral fellow, a visiting scholar, and two graduate research assistants.

We are also pleased to report that our group has recently received two large grants, one from NCR/Wal-Mart and the other from Hewlett-Packard (HP). The NCR grant is an unprecedented 1.6 terabyte parallel database engine valued at approximately $7 million. This system will help augment the other storage and compute servers we are purchasing through the NSF grant. This state-of-the-art parallel server has 80 processors and 400 disk drives. The grant from HP allows us to have 20 workstations and one server to use as desktops for the researchers in this project.

We have received an National Science Foundation ITR grant (joint with Ohio State University, who is a subcontractor) for studying multiple time scale traffic control aspects of the network research component. We have received a gift grant from the Xerox Foundation for studying the relationship between quality of service, network security, and fault-tolerance.

We received funding from the Santa Fe Institute (SFI) and National Science Foundation for holding a joint SFI/NSF Workshop, “The Internet as a Large-Scale Complex System” (Chair: Kihong Park, co-Chair: Walter Willinger), March 29-31, 2001. More information is available at http://discuss.santafe.edu/internetcs. We have received an NSF Career award to study efficient I/O techniques for a broad range of database applications. The project leverages the equipment that is being acquired through the MSI project.

2.2 Components and Materials Required (Inputs)

We have purchased and installed nearly $300,000 worth of equipment to support this project. The equipment can be broken down into four categories: (1) networking (four Cisco routers, 12 Intel workstations, two Sun workstations), (2) databases and storage (a Sun E450 server with A1000 disk array, an Intel-based storage server engine, three multimedia proxy servers, and two DVD storage carousels), (3) image compression (high resolution display), and (4) applications (MPEG-2 encoders, video camera for data input, three portable computers for distance learning, and three videoconferencing stations for connection with HBCU partners).

2.3 Indications of Success (Outputs)

Research subgroups have been established to address each of the principle areas of research. Because the overall goal of our work is the development of a unified infrastructure that encompasses many facets of multimedia database functionality, a research platform was created as a framework for integrating research results from all subgroups. Effective group interaction has been critical for establishing interfaces between the components of the infrastructure, and the successful implementation of this platform has provided a mechanism for integrating new techniques with minimal effort. This efficient method for testing and evaluating new ideas has led to rapid progress in every research area: the physical

storage management of multimedia data, quality-of-service (QoS)-based video access and retrieval, novel techniques for video shot segmentation based on feature extraction, multidimensional indexing structures to support queries against high dimensional feature vectors, and novel buffering techniques to support continuous quality-of- presentation (QoP) video streaming..

3 Evaluation

3.1 Degree of Success

We have accomplished our research objectives for the past year in each of the principle areas, and we plan further advancements in the development of components for our multimedia database infrastructure based on these successes. Specifically, in the area of storage management, we have designed and implemented several novel data placement, migration, pre-fetching, caching and scheduling schemes for effective video retrieval which satisfies QoS requirements. We have developed a proxy-based architecture that performs numerous real-time resource management functions, including bandwidth allocation, synchronization and transcoding of multimedia streams. Security and controlled access is handled through blocking or low resolution delivery, where users' views of the multimedia material is based on their security level. We have developed important new image processing techniques based on extracted visual features; these techniques have been incorporated into the image and semantic-based video shot segmentation methods used for meta-data indexing. We have made significant improvements to video query and browsing by applying search tree frameworks to multidimensional index structures for similarity searches against feature vectors. We have defined a high-level multimedia representation and MPEG-7 compliant description schemes (in XML format) to support multi-level video modeling and semantic video classification. We have designed an aggressive pre-fetching technique to support media streaming, and are investigating the effect of including the streaming operation in the query management pipeline. Continued research in these areas will result in a significant increase in the functionality and value of the MSI infrastructure.

3.2 Unmet Goals

No unmet goals have been identified.

3.3 Outcome and Impact

Recent advances in high-performance networks and improvements in computer hardware have created the potential for large-scale, distributed multimedia applications that can radically improve current real-world systems for distance learning, telemedicine, and e-commerce. The issues addressed by MSI research are the infrastructure components that make such applications possible. Traditional database methodology cannot be applied to multimedia databases, hence significant research, investigation and evaluation are required to identify innovative approaches for the storage, indexing, query, access, retrieval and presentation of video data that preserve quality-of-service and quality-of-presentation. We have made important progress in each of these research areas, developing the components for a basic infrastructure that supports not only our research platform "application" (for implementing and testing novel techniques), but also applications targeted for real-world use in distance learning and telemedicine.

4 Immediate Impact

4.1 Students

Mustafa Hammad (Fellowship, Ph.D. student)	Ihab Ilyas (RA, Ph.D. student)
Abdulmounaam Rezgui (RA, Ph.D. student)	Evimaria Terzi (RA, Ph.D. student)
Husni Fahmi (RA, Ph.D. student)	Wasfi Al-Khatib (Ph.D. student, graduated December 2000)
Muddasar Latif (RA, M.S. student)	Basit Shafiq (RA, M.S. student
James Joshi (RA, Ph.D. student)	Rahul Chari (RA, Ph.D. student)
Y. Xia (M. S. student)	Jiangtao Li (RA, Ph.D. student
Shaogang Chen (Ph. D., graduated May 2000)	John Cruz (RA, Ph.D. student)

Huan Ren (RA, Ph.D. student)	Tsunyi Tuan (RA, Ph.D. student)
Tapan Karwa	M. Elfeky (RA, Ph.D. student)
Ronaldo Ferreira	Shan Lei
Wei Li	Spencer Stanton (undergraduate senior)

4.2 Research/Programming Staff

N. Hirschburger, X. Zhu, J. Fan, D. Whittinghill, A. C. Catlin, A. Zhang, M. Marzouk

4.3 Visiting Scholars

A. Vakali, S. Hacid, E. Bertino

4.4 Publications

Elisa Bertino, Ahmed K. Elmagarmid and Mohand-Saïd Hacid. Quality of Service in Multimedia Digital Libraries. *SIGMOD Record*, 30(1), March 2001.

J. Fan and A.K. Elmagarmid, A Semi-automatic Semantic Algorithm for Video Object Extraction and Temporal Tracking, *Signal Processing: Image Communication*, 2001 (to appear).

J. Fan, W. G. Aref, M.-S. Hacid, and A. K. Elmagarmid, A Novel Color Edge Detection Technique, To appear, *Pattern Recognition Letters*, 2001.

K. Park, M. Sitharam, and S. Chen. Quality of service provision in noncooperative networks with diverse user requirements. *Decision Support Systems*, Special Issue on Information and Computation Economies, 28:101-122, 2000.

S. Dagtas, W. Al-Khatib, R. L. Kashyap, and A. Ghafoor. Motion Modeling and Indexing of Video Data. *IEEE Transactions on Image Processing*, January 2000.

W. Aref, I. Kamel, and S. Ghandeharizadeh. Disk Scheduling in Video Editing Systems. *IEEE Transactions on Knowledge and Data Engineering*, 2000 (to appear).

Jianping Fan, David Yau, Walid G. Aref, Ahmed K. Elmagarmid. Automatic Image Segmentation by Integrating Color Edge Extraction and Seeded Region Growing. *IEEE Transactions on Image Processing*, 2001 (to appear).

Jianping Fan, Walid G. Aref, and A. Rezgui. Adaptive Motion-Compensated Video Coding Scheme Towards Content-Based Bit Rate Allocation. *Journal of Electronic Imaging*, Vol. 9, No. 4, pp.521-533, 2000.

J. Fan, W. G. Aref, X. Zhu, I. Ilyas, M-S. Hacid, A. K. Elmagarmid. MultiView: Multi-level Video Content Representation and Retrieval. *Journal of Electronic Imaging*, October, 2001 (to appear).

K. Park and W. Willinger (eds.). *Self-Similar Network Traffic and Performance Evaluation*. Wiley-Interscience, 2000. Also includes Chapter 1: An overview (Park and Willinger) and Chapter 21: Future directions and open problems in performance evaluation and control of self-similar network traffic (Park).

H. Ren and K. Park. Toward a theory of differentiated services. In *Proceedings of the IEEE/IFIP International Workshop on Quality of Service*, 2000 (to appear).

H. Ren and K. Park. Efficient shaping of user-specified QoS using aggregate-flow control. To appear in *Proceedings of the International Workshop on Quality of Future Internet Services*, Lecture Notes in Computer Science, 2000.

T. Tuan and K. Park. Multiple time scale redundancy control for QoS-sensitive transport of real-time traffic. In *Proceedings of IEEE INFOCOM '00*, 2000.

T. Tuan and K. Park. Congestion Control for self-similar network traffic. Chapter 18 of *Self-Similar Network Traffic and Performance Evaluation*, Wiley-Interscience, 2000.

Husni Fahmi, Walid G. Aref, Muddasar Latif, and Arif Ghafoor. Distributed Framework for Real-Time Multimedia Object Communication. *Proceedings of the 3rd IEEE International Symposium on Object-Oriented Real-Time Distributed Computing*, pp. 252-259, Newport Beach, CA, March 2000.

5 Project Outcome

We have designed and implemented a web-based video retrieval system for medical education (EduMed) based on the distributed, multimedia database infrastructure developed through MSI research. EduMed is a significant and commercially viable system that offers valuable services for medical faculty and students, supporting rapid classroom access to interactive multimedia medical materials for teaching and learning. User-based security and access management allow faculty to prepare and deliver instruction through video query, retrieval, and presentation, while allowing students to access faculty-created video collections for research, assignments, and exams. The EduMed system provides end-to-end services for the analysis, storage, access, and retrieval of an indexed medical video library. The specific MSI-developed research components incorporated in the EduMed infrastructure are (1) processing for video indexing, using annotation, audio and content-based video shot analysis, (2) storage for video and meta-data with mechanisms to preserve quality-of-service (QoS) for video retrieval, (3) secure access, including protection for video ownership and patient privacy, (4) distributed multimedia databases, including a remote "archive" for raw video and indexing meta-data, as well as multiple local "warehouse" databases for user-built video shots collections resulting from queries against the archive, (5) video-enhanced query processing techniques, and (6) video streaming with continuous content flow mechanisms to preserve quality-of-presentation (QoP). EduMed was developed in partnership with researchers, educators, and industry service and product developers.

5.1 Databases

We have been able to build a truly multidisciplinary group among medical practitioners, medical educators, software developers, software vendors, and computer science researchers. As a result of this project we were able to receive additional funding from the State of Indiana to establish an incubator for the development of ideas resulting from this research. NCR and Wal-Mart jointly donated a large NCR Teradata system. A patent application is already on the way. We have also been able to participate in research in other areas of computer science, such as image processing, that enhance the efficacy and usefulness of our work.

5.2 Networking

The NSF RI grant has been instrumental in facilitating implementation, testing, and benchmarking of the QoS provisioning architecture developed by co-PI Park and his co-workers at the Network Systems Lab. This was achieved by building a dedicated IP-over-SONET backbone network comprised of Cisco 7206 VXR routers and collaborative support from Cisco Systems (technical contact: Fred Baker), which allows modification of the Cisco router operating system IOS to implement the QoS switching algorithms developed at the lab. The in-house QoS work is further leveraged in the Internet2 QoS Backbone (QBone) project where Park is a member of the QBone Architecture Design Team and QoS Working Group.

5.3 Security

Related to the network QoS project, the expansion of QoS provisioning to encompass security and fault-tolerance issues—a key challenge of the integrated QoS management objective of MSI and its scope—was achieved with assistance from the NSF RI grant. In particular, a new approach to distributed denial-of-service (DDoS) attack prevention called route-based distributed packet filtering was developed by co-PI Park and his team that allows scalable and proactive prevention of DDoS attacks on the global Internet. This research will be further funded by a new DARPA grant (PI: Park) from the FTN program, and expanded collaboration with Cisco Systems is being explored for possible IETF standardization and BGP-based implementation in inter-domain routing.