MSI: A Research Infrastructure for Integrated Quality of Service Management in Multimedia Computing Environments

W. Aref, A.K. Elmagarmid, S. Fahmy, A. Ghafoor, T. Korb, K. Park, S. Prabhakar, E. Spafford, D. Yau
Purdue University, West Lafayette, IN, 47907

1. Introduction

Purdue University is a land grant institution located in West Lafayette Indiana. The departments of computer sciences and electrical and computer engineering boast a long tradition of excellence in Computer Science. Among other distinctions, the computer sciences department was the first institution to offer a doctoral degree in computer science. Purdue University is also host to the Center for Education and Research in Information Assurance and Security (CERIAS), which was named a national center of excellence in information security education.
This is the first year of the RI award for the MSI project.

The main project web page is http://www.cs.purdue.edu/msi
Other related pages include:
http://www.cs.purdue.edu/icds (Indiana Center for Database Systems)
http://shay.purdue.edu/~dmultlab (Distributed Multimedia Systems Lab home page)
http://www.cs.purdue.edu/nsl (Network Systems Lab home page)
http://www.cerias.purdue.edu/ (CERIAS home page)

2. Goals

The overall objective of the Purdue Multimedia Support Infrastructure (MSI) project is to address research challenges for developing a prototype system that will integrate key information technologies – database and storage management, networking, and security – in order to support a comprehensive end-to-end Quality-of-Service (Qos) management framework for distributed multimedia document applications. The term QoS deals with the provision for user specified quality requirements, including security, reliability, precision, etc. The heterogeneous notion of the quality embodied by various technologies of MSI raises several challenges issues in capturing, storage, delivery and presentation of multimedia information with a view to optimize the overall performance. By combining theory with extensive experimentation, the proposed infrastructure will provide a generic framework that is expected to support a variety of distributed multimedia applications requiring QoS guarantees.

3. Accomplishments

The overall research activities for the MSI project fall into three technology areas, namely: multimedia database and storage management, networking and security. We now briefly describe activities and research outcomes in these areas:

3.1 Multimedia Database and Storage Management Subsystem

A novel, multi-layered, real-time, distributed multimedia database system (RTDMDS) architecture with Quality of Service (QoS) guarantees has been proposed by the investigators. The system allows distributed users to author, store, query, and retrieve multimedia documents over a broadband network. The system uses a Petri-net based document model for modeling of document meta-data, querying, and communication of multimedia documents. This model is an enhanced version of our earlier proposed model for multimedia document composition. The enhanced model allows specification of user level Quality of Presentation (QoP) parameters and content information of individual media in a document. The model has several desired features that facilitate content based-based retrieval of documents, determination of system-wide resource requirements to guarantee QoS, and implementation of stream scheduling mechanisms for storage and server subsystems.

The overall architectural framework for RTDMDS is based on some of the most advanced techniques and models developed for multimedia document management and content-based retrieval of multimedia information. In particular, several novel techniques have been proposed by the investigators for the QoP-based specification of complex multimedia objects, video data modeling for efficient searching and querying of multimedia data and access control mechanisms for multimedia database systems. These techniques have the potentials of significantly advancing the Web-based multimedia information technology.

We are also pursuing research issues related to the design and development of a video database system. The objective is to allow content-based retrieval and indexing of large-scale video archives, which in turn will be a key component of a multimedia document management system.

For storage management, we are focusing on two areas of research, namely, storage management of video data and placement of complex multimedia documents. For video data, we are analyzing several techniques for real-time disk scheduling. For multimedia documents, we have designed several novel data placement and scheduling schemes that are currently being tested through extensive simulation. These techniques are also being implemented on a Sun E450 server and the Sun A1000 RAID array.

3.2 Networking Subsystem

We have set forth a new unified theory of differentiated services that is implementable on IP networks. It improves on our earlier work on QoS scheduling, facilitating a theory of aggregate-flow QoS control.
We have developed a framework for reasoning about the "goodness" of various differentiated services architectures that considers both the per-hop behavior (PHB) and edge control parts when provisioning user-specified QoS-sensitive network services. We have benchmarked the architecture over QSim, an ns-based WAN QoS simulation environment, and shown that optimal aggregate-flow per-hop behavior can export efficient differentiated services commensurate with user requirements.

As a starting point to implementing a full-fledged version of our optimal aggregate-flow classifier inside Cisco 7206 VXR routers, we have implemented end-to-end signaling controls, both for current IETF Diff-Serv specifications (i.e., Assured Service) and IETF integrated services (Int-Serv) using RSVP. Our end-to-end signaling platform can be adapted, with modular changes, to the optimal aggregate-flow per-hop control architecture advanced in our theoretical and simulation work. We have tested the real-time scheduling and switching capabilities of a network of Cisco 7206 VXR routers connected as an IP-over-SONET internetwork with 100 Mbps network interfaces tapping end station PCs -- serving both as QoS end points and traffic generators -- and ascertained their operating dimensions and properties with respect to yielding predictable performance. We are in the process of implementing the first version of optimal aggregate-flow PHB inside Cisco 7206 VXR routers, which will then be tested with respect to its QoS provisioning.

We have developed and implemented end-to-end QoS amplification techniques using adaptive packet-level forward error correction and multiple time scale traffic control. We have advanced the multiple time scale (MTS) traffic control framework that exploits large time scale correlation structure present in bursty—in particular, self-similar—network traffic for traffic control purposes. MTS traffic control facilitates a much-needed measure of proactivity, which, in turn, facilitates scalable end-to-end traffic and QoS control. MTS traffic control has been implemented both for QoS (a multiple time scale extension of AFEC called AFEC-MT), and throughput maximization for TCP (MTS extension called TCP-MT over Reno, Tahoe, and Vegas) and rate-based control for ATM. Multiple time scale redundancy control (i.e., AFEC-MT) has been benchmarked over UDP/IP internetworks that transport real-time MPEG I & II video and audio. The systems run both for UNIX and Windows NT.

3.3 Security Subsystem

For the security subsystem, we are focusing our efforts to build security-based access control mechanisms for video and multimedia databases. For video databases, we are currently developing an access control system on top of a video database system. At the higher levels, we are developing an access control model that specifies the users' credentials and qualifications as well as the content description of the underlying video. On the other hand, at the lower levels, we are building a toolbox for extracting desired features from the underlying video streams. For multimedia document system, we are developing a security framework to allow integration of heterogeneous access control policies in a distributed environment.

For developing a secure network, we have proposed a framework for probabilistic packet marking for distributed denial-of-service attack (DoS) prevention which puts forth an effective technique for dealing with an important security and QoS threat. Instead of logging path information at routers which makes traceback of DoS attack overhead-prone and ineffective in high-speed networks where routers are expected to switch at Gbps or higher rates, the complexity is pushed to the edge following the end-to-end paradigm, and "logging" is carried out by probabilistic sampling of path information in an end-to-end manner. This method can be easily implemented in IP routers. Our work provides a comprehensive analysis of the power of probabilistic

packet marking (PPM) with respect to both single-source and distributed DoS. We use a 2-player adversarial framework to analyze the minimax and maximin optimal strategies of the attacker and victim, and have concluded that PPM is effective at preventing DoS attacks. Using numerical results, it has been shown that the outcome of optimal strategies can be achieved when using typical Internet based configurations including hop count, topology, spoofing, and marking probability.

4. Impact and Indications of Success

So far the project goals for the first year have been successfully accomplished. For example, we have been extremely successful in leveraging the NSF CISE RI funds in several significant and noticeable ways. In particular, we 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.

In addition, we have established research collaboration with Telcordia and Siemens Corporation. We have also received substantial internal research support from the Center for Education and Research in Information Assurance and Security (CERIAS).

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 be used as a backing store and will help augment the other storage and compute servers we are purchasing through the NSF grant.

Post-Doctoral Researchers
Four post-doctoral researchers are involved in research related to the MSI project:
J. Fan, H. Lee, and M. Hacid, and W. McIver. In addition there are three programmers working on the project: N. Hirschberger, M. Marzouk, and D. Whitingfield.

Students
16 graduate students are conducting research related to the equipment that has been purchased and will be purchased through the RI grant. They are: W. Al-Khatib, R. Chari, S. Chen, J. Cruz, H. Fahmi, M. Hammad, I. Ihab, J. Joshi, M. Latif, J. Li, H. Ren, A. Rezgui, B. Shafiq, E. Terzi, T. Tuan, and Y. Xia. Of these, two are women (E. Terzi and Y. Xia).

Project References

W. Al-Khatib and A. Ghafoor. An Approach for Video Data Modeling and Query Processing. *Proceedings of ACM Multimedia'99*, November 1999. Orlando.

S. Chen and K. Park. An architecture for noncooperative QoS provision in many-switch systems. In *Proceedings of IEEE INFOCOM '99*, pp. 864-872, 1999.

S. Dagtas and A. Ghafoor. Indexing and Retrieval of Video Based Spatial Relation Sequences. *Proceedings of ACM Multimedia '99*, November 1999. Orlando.

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.

Jianping Fan, Walid G. Aref, and A. Rezgui. Adaptive Motion-Compensated Video Coding Scheme Towards Content-Based Bit Rate Allocation. *Journal of Electronic Imaging*. To appear, 2000.

Husni Fahmi, Walid G. Aref, Mudassir 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.

Jianping Fan, Walid G. Aref, and A. Rezgui. Automatic Object-Oriented Segmentation and Conceptual Modeling Algorithm towards Content-Based Representation. *Proceedings of the International Conference on Visual Communication and Image Processing*, Perth, Australia, June 2000 (to Appear).

F. Kokkoras, H. Jiang, I. Vlahavas, A. K. Elmagarmid, E. N. Houstis, and W. Aref. Smart Video Text: An Intelligent Video Database System. *Journal of ACM Multimedia Systems*. To appear, 2000.

P. Liu and A. Ghafoor (eds.). *Multimedia Document Systems: A Perspective*. Published by Kluwer Academic Publishers, 1999.

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.

K. Park and W. Willinger (eds.). *Self-Similar Network Traffic and Performance Evaluation*. Wiley-Interscience, 2000.

H. Ren and K. Park. Toward a theory of differentiated services. In *Proceedings of the IEEE/IFIP International Workshop on Quality of Service*, pp. 211-220, 2000.

T. Tuan and K. Park. Multiple time scale congestion control for self-similar network traffic. *Performance Evaluation*, 36:359-386, 1999.

T. Tuan and K. Park. Performance evaluation of multiple time scale TCP under self-similar traffic conditions. *To appear in ACM Transactions on Modeling and Computer Simulation*, 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.