

Annual Report for Period:04/2001 - 04/2002**Submitted on:** 05/29/2002**Principal Investigator:** Elmagarmid, Ahmed K.**Award ID:** 9972883**Organization:** Purdue Research Foundation**Title:**

CISE Research Infrastructure: MSI: A Research Computing Environments

Infrastructure for Integrated Quality of Service Management in Multimedia

Project Participants**Senior Personnel****Name:** Elmagarmid, Ahmed**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Korb, John**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Ghafoor, Arif**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Spafford, Eugene**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Park, Kihong**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Aref, Walid**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Fahmy, Sonia**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Prabhakar, Sunil**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Yau, David**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Blair, Shane**Worked for more than 160 Hours:** Yes**Contribution to Project:**

Name: Delp, Ed

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Szpankowski, Wojtek

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Catlin, Ann Christine

Worked for more than 160 Hours: Yes

Contribution to Project:

Post-doc

Name: Fan, Jianping

Worked for more than 160 Hours: Yes

Contribution to Project:

Jianping Fan (Ph.D., Post-Doctorate) is pursuing his research in providing low-level support in video databases including, feature extraction, spatial and temporal segmentation, object detection and recognition, object motion tracking, and multimedia database indexing.

Name: Hacid, M.

Worked for more than 160 Hours: Yes

Contribution to Project:

M. Hacid (Ph.D., Post-Doctorate) recently joined the MSI project. His main role will be to design and oversee the implementation of the QoS/QoS Interface.

Name: Lee, Heejo

Worked for more than 160 Hours: Yes

Contribution to Project:

Heejo Lee has begun preliminary work on ascertaining Cisco's network security framework for subsequent adaptation to the AdSec network security architecture advanced in the project. Lee and Park (co-PI) have written a paper on 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.

Name: Vakali, Athena

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Bahk, Saewoong

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Kim, Jae-Kwan

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Kim, Jisoo

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Kim, Jae-Young

Worked for more than 160 Hours: Yes

Contribution to Project:

A new postdoc in the Network Systems Lab who is handling the measurement and workload generation aspects of the Q-Bahn architecture.

Name: Zhu, Xing Quan

Worked for more than 160 Hours: Yes

Contribution to Project:

Graduate Student

Name: Hammad, Mustafa

Worked for more than 160 Hours: Yes

Contribution to Project:

Mustafa Hammad (Purdue Fellowship, Ph.D. student) is pursuing his research in developing access control mechanisms for video database systems both at the frame and object levels.

Name: Ilyas, Ihab

Worked for more than 160 Hours: Yes

Contribution to Project:

Ihab Ilyas (RA, Ph.D. student) is pursuing his research in developing multimedia indexing techniques to support fast content-based retrieval of multimedia objects in the context of video databases.

Name: Rezgui, Abdoulmounaam

Worked for more than 160 Hours: Yes

Contribution to Project:

Abdoulmounaam Rezgui (RA, Ph.D. student) is pursuing his research in developing scalable multimedia servers that support real-time streaming of video data, quality of service, and quality of presentation.

Name: Terzi, Evimaria

Worked for more than 160 Hours: Yes

Contribution to Project:

Evimaria Terzi (RA, Ph.D. student) will begin as an RA on the multimedia database project in August 2000.

Name: Fahmi, Husni

Worked for more than 160 Hours: Yes

Contribution to Project:

Husni Fahmi (defended PhD thesis in April 2002). Husni Fahmi's research is related to the development of an overall framework for a proxy-based real-time distributed multimedia document system (RTDMDS) to allow searching and retrieval of distributed multimedia documents.

Name: Al-Khatib, Wasfi

Worked for more than 160 Hours: Yes

Contribution to Project:

Wasfi Al-Khatib (RA, Ph.D. student) completed his doctoral studies and has been trained in the area of multimedia databases. In particular, he has developed models for indexing and content-based retrieval of video data. He has published papers in several conferences and journals.

Name: Latif, Muddasar

Worked for more than 160 Hours: Yes

Contribution to Project:

Muddasar Latif (RA, M.S. student) is pursuing his MS thesis in multimedia document delivery over the Web. He is involved in the development of different modules of RTDMDS.

Name: Shafiq, Basit

Worked for more than 160 Hours: Yes

Contribution to Project:

Basit Shafiq (RA, Ph.D. student) is pursuing his doctoral thesis in multimedia document modeling and search mechanism. He has completed the development of several modules for RTDMDS. Currently, he is pursuing his work on the synchronization and integration aspects of heterogeneous multimedia databases.

Name: Joshi, James

Worked for more than 160 Hours: Yes

Contribution to Project:

James Joshi (RA, Ph.D. student) is pursuing his research in the area of multimedia database security. Currently he is developing a framework that will allow integration of heterogeneous security and access control mechanisms in a distributed multimedia document environment. The framework will integrate both Role-Based Access Control (RBAC) models and Task-Based Access Control (TBAC) models.

Name: Chari, Rahul

Worked for more than 160 Hours: Yes

Contribution to Project:

Rahul Chari (RA, Ph.D. student) is developing a simulator for the storage management aspects of MSI, and developing and testing alternative schemes.

Name: Xia, Y.

Worked for more than 160 Hours: Yes

Contribution to Project:

Y. Xia is a graduate student working on an independent study project in the storage area.

Name: Li, Jiangtao

Worked for more than 160 Hours: Yes

Contribution to Project:

Jiangtao Li (RA, Ph.D. student) is implementing different storage management schemes for their testing and evaluation.

Name: Chen, Shaogang

Worked for more than 160 Hours: Yes

Contribution to Project:

Shaogang Chen (RA, Ph.D. student) has graduated May 2000, with his Ph.D. dissertation 'Stratified Best-effort QoS Provisioning in Noncooperative Networks' advancing the state-of-the-art in the area. Chen has joined CPlane, a start-up company that focuses on QoS and network management switching products.

Name: Cruz, John

Worked for more than 160 Hours: Yes

Contribution to Project:

John Cruz (RA, Ph.D. student) has performed end system scheduling work for end-to-end QoS using, in part, the Purdue Infobahn network QoS infrastructure.

Name: Ren, Huan

Worked for more than 160 Hours: Yes

Contribution to Project:

Huan Ren (RA, Ph.D. student) has published two recent papers on differentiated services QoS with a third paper presently in submission.

Name: Tuan, Tsunyi

Worked for more than 160 Hours: Yes

Contribution to Project:

Tsunyi Tuan (RA, Ph.D. student) has published one journal paper, one book chapter, and two conference papers based on traffic control work related to the project.

Name: Ferreira, Ronaldo

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Lei, Shan

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Li, Wei

Worked for more than 160 Hours: Yes

Contribution to Project:

Name: Ren, Zhong

Worked for more than 160 Hours: Yes

Contribution to Project:

A research assistant in the Network Systems Lab who has worked on the QoS benchmarking and testing aspect of the project.

Undergraduate Student

Name: Stanton, Spencer

Worked for more than 160 Hours: Yes

Contribution to Project:

Spencer Stanton (senior) did an independent study with Park (co-PI) in the spring semester where he assisted graduate students with implementing the MPEG audio part of AFEC for Optibase real-time compression board under Windows NT.

Technician, Programmer

Name: Marzouk, Mirette

Worked for more than 160 Hours: Yes

Contribution to Project:

Mirette Marzouk is a staff programmer involved with general application support, Web page design, and software development.

Name: Hirschberg, Nicholas

Worked for more than 160 Hours: Yes

Contribution to Project:

Nicholas Hirschberg maintains commercial databases (Oracle and Teradata) for the project and provides general support of the Web server.

Name: Zhang, X.

Worked for more than 160 Hours: Yes

Contribution to Project:

X. Zhang is a visiting scholar who started in September 2000. She will serve as a programmer working on object identification for the video database system.

Name: Whittinghill, David

Worked for more than 160 Hours: Yes

Contribution to Project:

Other Participant

Research Experience for Undergraduates

Organizational Partners

Cisco Systems Inc

Clarian Health

Indiana University Medical School

Florida A&M University

We have purchased and delivered to Florida A&M University a videoconferencing system to help them participate in our research seminars.

Indiana Health Industry Forum

Med Institute

Methodist Hospital of Indiana Inc

Micro Data Base Systems Inc

Ohio State University

Siemens Corporate Research

Spelman College

We have purchased and delivered to Spelman College a videoconferencing system to help them participate in our research seminars.

University of Delaware

University of Notre Dame

Internet2

Santa Fe Institute

Tuskegee University

We have purchased and delivered to Tuskegee University a videoconferencing system to help them participate in our research seminars.

Xerox Corporation

Other Collaborators or Contacts

Fred Baker (Cisco)

E. Bertino (University of Milan)

Jennifer Hou (Ohio State University)

Andrea Lawrence (Spelman College)

Peiya Liu (Siemens Corporation)

Jon Reid (MDBS)

V. Verykios (Drexel University)

Deidre Williams (Florida A&M University)

Ashfaq Khokhar (University of Delaware)

Ben Teitelbaum (Internet2)

Suzanne Dulle (Santa Fe Institute)

Hira Narang (Tuskegee University)
 Paul Barford (Univ. of Wisconsin)
 Azer Bestavros (Boston Univ.)
 Walter Willinger (AT&T Research)

Activities and Findings

Research and Education Activities: (See PDF version submitted by PI at the end of the report)

See attached file.

Findings: (See PDF version submitted by PI at the end of the report)

See attached file.

Training and Development:

The advanced multimedia course 'Multimedia Information Systems' taught by Ghafoor (co-PI) in Fall 1999 used the multimedia database servers (video data server) and exposed graduate students to state-of-the-art multimedia database and document management techniques.

The exploratory advanced networking course 'Multimedia and High-Speed Networks' taught by Park (co-PI) in Spring 2000 used the evolving Purdue Infobahn QoS Testbed, and exposed graduate students to state-of-the-art network QoS techniques and implementations. An extension of this course has recently been approved by the CS faculty as a regular course (CS 637).

Leveraging the QoS and traffic control work carried out in the Network Systems Lab, co-PI Kihong Park has forged a collaborative relationship with Internet2 to perform wide area network (WAN) QoS amplification experiments over Abilene/QBone, the Gbps backbone of Internet2.

The Advanced Topics in Database Systems course taught by Walid Aref (member of the MSI Senior Personnel) for the Computer Science Department in Fall 1999 and Fall 2001 covers new trends in database systems, including multimedia database server design issues, advanced query processing and query optimization techniques, generalized search trees (GIST), and many other topics profoundly related to research on the multimedia database server for the MSI initiative. Graduate students and research staff presented lectures in state-of-the-art database technology, describing design, techniques, algorithms, implementation, and performance evaluation issues related to the database server, as well as demonstrations of the operational system.

Outreach Activities:

We have engaged in several outreach activities.

ò Involved four Historically Black Colleges and Universities (Fisk, Florida A&M, Spelman, and Tuskegee) in our research program. We have invited them to workshops and presentations and provided videoconferencing equipment so they can attend seminars remotely. We arranged summer internships for students at Telcordia Technologies. We will continue to interact with the students and encourage them to apply to the Masters and Ph.D. programs at Purdue.

. Kihong Park has interacted with members of the Santa Fe Institute in the capacity of Fellow-at-Large. Park also organized a workshop at the Santa Fe Institute (in February 2001) called 'The Internet as a Complex System', which looked at scaling, performance, and control issues related to the Next Generation Internet. It was a by-invitation only meeting, and selected contributions will be published by Oxford University Press as part of the SFI book series.

. Park has given invited presentations at SFI to its Business Network members, which includes CTOs and managers from a range of institutions and organization in the US and abroad. Park has also given presentations in the Dean's Honors Seminar at Purdue University conducted by the Associate Dean, which has exposed some of the most talented undergraduate students in the School of Science to research conducted in the Network Systems Lab and the MSI project.

. Park will also give a tutorial entitled 'Internet Traffic Modeling and its Implications to Network Performance and Control' at Performance 2002, in Rome, Italy, which will expose the broader networking audience, including professionals and engineers, to the latest research results in traffic and workload modeling (an important component of the MSI project) which has taken place over the past decade.

Journal Publications

J. Cruz and K. Park, "Towards performance-driven system support for distributed computing in clustered environments", Journal of Parallel and Distributed Computing, p. 132, vol. 59, (1999). Published

- 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*, p. 101, vol. 28, (2000). Published
- T. Tuan and K. Park, "Multiple time scale congestion control for self-similar network traffic", *Performance Evaluation*, p. 359, vol. 36, (1999). Published
- 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*, p. , vol. , (). Accepted
- S. Dagtas, W. Al-Khatib, R. L. Kashyap, and A. Ghafoor, "Motion Modeling and Indexing of Video Data", *IEEE Transactions on Image Processing*, p. 1, vol. , (2000). Published
- W. Aref, I. Kamel, and S. Ghandeharizadeh, "Disk Scheduling in Video Editing Systems", *IEEE Transactions on Knowledge and Data Engineering*, p. , vol. , (2000). Accepted
- Jianping Fan, Walid G. Aref, and A. Rezgui, "Adaptive Motion-Compensated Video Coding Scheme Towards Content-Based Bit Rate Allocation", *Journal of Electronic Imaging*, p. x, vol. , (2000). Accepted
- David K. Y. Yau and Xiangjing Chen, "Resource Management on Software-Programmable Router Operating Systems", *IEEE Journal on Selected Areas in Communications*, p. x, vol. 19, (2001). Published
- David K. Y. Yau, "Performance Evaluation of CPU Isolation Mechanisms in a Multimedia OS Kernel", *Proceedings of the IS&T/SPIE Multimedia Computing and Networking*, p. , vol. , (2001). Accepted
- David K. Y. Yau and Bharat Bhargava, "Heterogeneous CPU Services using Differentiated Admission Control", *Kluwer Multimedia Tools and Applications Journal*, p. , vol. , (). Submitted
- Shivkumar Kalyanaraman, Raj Jain, Sonia Fahmy, Rohit Goyal, and Bobby Vandalore, "The ERICA Switch Algorithm for ABR Traffic Management in ATM Networks", *IEEE/ACM Transactions on Networking*, p. 87, vol. 8, (2000). Published
- Bobby Vandalore, Sonia Fahmy, Raj Jain, Rohit Goyal, and Mukul Goyal, "General Weighted Fairness and its Support in Explicit Rate Switch Algorithms", *Computer Communications Journal*, p. 149, vol. 23, (2000). Published
- Bobby Vandalore, Wu-chi Feng, Raj Jain, and Sonia Fahmy, "A Survey of Application Layer Techniques for Adaptive Streaming of Multimedia", *Journal of Real Time Imaging*, p. 221, vol. 7, (2001). Published
- M. Frantzen, F. Kerschbaum, E. Schultz, and S. Fahmy, "A Framework for Understanding Vulnerabilities in Firewalls Using a Dataflow Model of Firewall Internals", *Journal of Computers and Security*, p. 263, vol. 20, (2001). Published
- Sonia Fahmy, Raj Jain, Rohit Goyal, and Bobby Vandalore, "Fair Flow Control for ATM-ABR Multipoint Connections", *Journal of Computer Communications*, p. 741, vol. 25, (2002). Published
- Sonia Fahmy, Raj Jain, Rohit Goyal and Bobby Vandalore, "On Determining the Fair Bandwidth Share for ABR Connections in ATM Networks", *Journal of High Speed Networks*, p. , vol. , (2002). Accepted
- J. Joshi, A. Ghafoor, W.G. Aref, E. H. Spafford, "Digital Government Security Infrastructure Design Challenges", *IEEE Computer*, p. 66, vol. 34, (2001). Published
- J. Joshi, W. G. Aref, A. Ghafoor and E. H. Spafford, "Security models for web-based applications", *Communications of the ACM*, p. 38, vol. 44, (2001). Published
- J. Joshi, A. Ghafoor, "A Petri-Net Based Multilevel Security Specification Mechanism for Multimedia Documents in a Multidomain Environment", *The Second Annual Systems Security Engineering Conference*, p. x, vol. , (2001). Published

- J. Joshi, A. Ghafoor, "A Petri-Net Based Multilevel Security Specification Model for Multimedia Documents", IEEE International Conference on Multimedia and Expo, p. x, vol. , (2000). Published
- Z. Ali, A. Ghafoor and C.S.G. Lee, "Media Synchronization in Multimedia Web Using a Neuro-Fuzzy Framework", IEEE Journal of Selected Areas in Communications, p. 168, vol. 18, (2000). Published
- Z. Ali and A. Ghafoor, "Synchronized Delivery of Multimedia Information over ATM Networks", Communications of the ACM, p. 239, vol. 43, (2000). Published
- H. Fahmi, M. Latif, A. Ghafoor, P. Liu, and L. Hsu, "Performance Evaluation of an Interactive Web-Based Multimedia Document", IEEE International Conference on Multimedia and Expo (ICME 2000), p. x, vol. , (2000). Published
- H. Fahmi, W. G. Aref, M. Latif, B. Shafiq, A. Ghafoor, "Distributed Framework for Real-Time Multimedia Object Communication", Journal of Computer System Science and Engineering, p. , vol. , (2001). Accepted
- J. Joshi, Z. K. Li, H. Fahmi, B. Shafiq, and A. Ghafoor, "A Model for Secure Multimedia Document Database Systems", IEEE Transactions on Multimedia, p. , vol. , (). Submitted
- K. Park and T. Tuan, "Performance evaluation of multiple time scale TCP under self-similar traffic conditions.", ACM Transactions on Modeling and Computer Simulation. SPecial Issue on Modeling and Communication Networks, p. 1, vol. 10, (2000). Published
- K. Park and W. Wang, "QoS-sensitive transport of real-time MPEG video using adaptive redundancy control.", Computer Communications, Special Issue on QoS-Sensitive Distributed Networks Systems and Applications, p. 78, vol. 24, (2001). Published
- Elisa Bertino, Ahmed K. Elmagarmid and Mohand-Sa'd Hacid, "Quality of Service in Multimedia Digital Libraries", SIGMOD Record, p. , vol. 30, (2001). Accepted
- Mohand-Sa'd Hacid, Farouk Toumani and Ahmed K. Elmagarmid., "Constraint-Based Approach to Semistructured Data.", Fundamenta Informaticae Journal, p. , vol. , (2001). Accepted
- Jianping Fan, David K.Y. Yau, Walid G. Aref and A. Rezgui, "Adaptive Motion-compensated Video Coding Scheme Towards Content-based Bit-rate Allocation", IS&T/SPIE Journal of Electronic Imaging, p. x, vol. 9, (2000). Published
- Jianping Fan, Mohand S. Hacid and Ahmed K. Elmagarmid., "Content-based video classification toward hierarchical representation, indexing and accessing.", Journal of Multimedia Tools and Applications, p. , vol. , (2001). Accepted
- A. Vakali, E. Terzi, and A. Elmagarmid., "Representation and Storage Modeling in Multimedia Systems", ', Journal of Applied Systems Studies, Special Issue on Distributed Multimedia Systems with Applications, p. , vol. 2, (2001). Accepted
- A. Vakali and E. Terzi, "Video Data Storage Policies : An Access Frequency Based Approach", Computers & Electrical Engineering Journal, Elsevier, p. , vol. , (2001). Accepted
- Jianping Fan, Walid G. Aref, Mohand S. Hacid and Ahmed K. Elmagarmid, "An improved isotropic color edge detection algorithm", Pattern Recognition Letters, p. 1419, vol. , (2001). Published
- A. Vakali and E. Terzi, "A Java-based model for I/O scheduling in Tertiary Storage Subsystems", International Journal of Computers and Applications, ACTA Press, p. , vol. , (2001). Accepted
- Rohit Goyal, Raj Jain, Sonia Fahmy, Bobby Vandalore, Mukul Goyal, Sastri Kota, Norm Butts, Kul Bhasin, and Thomas VonDeak, "Buffer Management and Rate Guarantees for TCP/IP over Satellite-ATM Networks", International Journal of Satellite Communications, p. 111, vol. 19, (2001). Published

- Sonia Fahmy, Raj Jain, Rohit Goyal, and Bobby Vandalore, "Design and simulation of ATM-ABR end system congestion control", Transactions of the Society for Computer Simulation, p. , vol. , (2002). Accepted
- J. Fan, D. K. Y. Yau, A. K. Elmagarmid, W. G. Aref, "Automatic Image Segmentation by Integrating Color Edge Detection and Seeded Region Growing", IEEE Transactions On Image Processing, p. 1454, vol. 10, (2001). Published
- J. Fan, W. G. Aref, A. K. Elmagarmid, M. S. Hacid, M. S. Marzouk, X. Zhu, "Multiview: Multi-level Video Content Representation and Retrieval", Journal of Electrical Imaging, p. 895, vol. 10, (2001). Published
- W. Aref, I. Kamel, and S. Ghandeharizadeh, "Disk Scheduling in video editing systems", IEEE Transactions on Knowledge and Data Engineering, p. 933, vol. 13, (2001). Published
- J. Fan, M. S. Hacid, A. K. Elmagarmid, "Model-based Video Classification for Hierarchical Video Access", Multimedia Tools and Applications, p. 89, vol. 17, (2002). Published
- J. Fan and A. K. Elmagarmid, "An Automatic Algorithm for Semantic Object Generation and Temporal Tracking", Signal Processing: Image Communication, p. , vol. 17, (2002). Accepted
- J. Fan and A. K. Elmagarmid, "Semi-automatic algorithm for video object extraction and temporal tracking", Signal Processing: Image Communication, p. , vol. 17, (2002). Accepted
- J. Fan, X. Zhu, M. S. Hacid, A. K. Elmagarmid, "Cluster-based indexing of video database: towards more efficient retrieval and browsing", IEEE Transactions on Image Processing, p. , vol. 11, (2002). Accepted
- J. Fan, X. Zhu, A. K. Elmagamid, and W. G. Aref, "MultiView: Hierarchical video classification and access system", IEEE Transactions on Multimedia, p. 895, vol. 4, (2002). Published
- W. Aref, M. Elfeky, and A. K. Elmagarmid, "Incremental, online and merge mining of partial periodic patterns in time-series databases", IEEE Transactions on Knowledge and Data Engineering, p. , vol. , (2002). Accepted
- X. Zhu, A. K. Elmagarmid, and A. C. Catlin, "InsightVideo: Toward hierarchical content organization for efficient browsing, summarization and retrieval", IEEE Transactions on Multimedia, p. , vol. , (2002). Submitted
- E. N. Houstis, S. Fleeter, J. R. Rice, and C. Zhou, "GasTurbnLab: A problem solving environment for gas turbine engine simulation", Journal of Comp. Engr. Sci., p. , vol. , (2001). Accepted
- E. Deelman, A. Dube, A. Hoisie, E. Houstis, Y. Luo, R. L. Oliver, D. Sundaram-Stukal, H. Wasserman, V. S. Adve, R. Bagrodia, J. C. Browne, O. Lubeck, J. R. Rice, P. J. Teller, and M. K. Vernon, "POEMS: End-to-end performance design of large parallel adaptive computational systems", IEEE Trans. Soft. Engr., p. 1027, vol. 26, (2000). Published
- N. Bassiliades, I. Vlahavas, and A. K. Elmagarmid, and E. N. Houstis, "InterBaseKB: Integrating a Knowledge Base System with a Multidatabase System for Data Warehousing", IEEE Trans. Of Knowledge and Data Engineering, p. , vol. , (2002). Accepted
- E. N. Houstis, A. C. Catlin, P. Tsompanopoulou, D. Gottfried, G. Balakrishnan, K. Su, and J. R. Rice, "GasTurbnLab: A Multidisciplinary Problem Solving Environment for Gas Turbine Engine Design on a Network of Non-Homogeneous", Journal of Computational and Applied Mathematics, p. , vol. , (2002). Accepted
- E. N. Houstis, A. Catlin, N. Ramakrishnan, N. Dhanjani, V. Verykios, and J. R. Rice, "MyPythia: A recommendation portal for scientific software and services", Concurrency: Practice and Experience, Special Issue on Scientific Portals, p. , vol. , (2002). Accepted
- A. Habib, M. Hefeeda, and B. Bhargava, "Detecting Service Violations and DoS Attacks", IEEE Network, p. , vol. , (2002). Submitted
- A. Habib, S. Fahmy, S. Avasarala, V. Prabhakar, and B. Bhargava, "On Detecting Service Violations and Bandwidth Theft in QoS Network Domains", Computer Communications, p. , vol. , (2002). Submitted

- B. Bhargava, "A Study of Communication Delays for Web Transactions", Cluster Computing Journal, p. 319, vol. 4, (2001). Published
- B. Bhargava and M. Annamalai, "A Communication Framework for Digital Libraries", Journal of Multimedia Systems and Applications, p. 205, vol. 10, (2000). Published
- H. Fahmi, M. Latif, S. Sedigh-Ali, A. Ghafoor, P. Liu, and L.H. Hsu, "Proxy Servers for Scalable Interactive Video Support", IEEE Computer, p. 54, vol. 34, (2001). Published
- Sunil Prabhakar, Y. Xia, D. Kalashnikov, W. Aref, and S. E. Hambrusch, "Query Indexing and Velocity Constrained Indexing: Scalable Techniques for Continuous Queries on Moving Objects", IEEE Transactions on Computers, Special Issue on DBMS and Mobile Computing, p. , vol. , (2002). Accepted

Books or Other One-time Publications

- P. Liu and A. Ghafoor, "Multimedia Document Systems: A Perspective", (1999). Book, Published
Editor(s): P. Liu and A. Ghafoor
Bibliography: Kluwer Academic Publishers
- S. Chen, "Stratified Best-effort QoS Provisioning in Noncooperative Networks", (2000). Thesis, Published
Bibliography: Purdue University
- S. Chen and K. Park, "An architecture for noncooperative QoS provision in many-switch systems", (1999). Proceedings, Published
Collection: INFOCOM '99
Bibliography: IEEE
- K. Park, "Future directions and open problems in performance evaluation and control of self-similar network traffic", (2000). Book Chapter, Published
Collection: Self-Similar Network Traffic and Performance Evaluation
Bibliography: Wiley-Interscience
- K. Park and W. Wang, "QoS-sensitive transport of real-time MPEG video using adaptive forward error correction", (1999). Proceedings, Published
Collection: Multimedia Systems '99
Bibliography: IEEE
- K. Park and W. Willinger, "Self-Similar Network Traffic and Performance Evaluation", (2000). Book, Published
Editor(s): K. Park and W. Willinger
Bibliography: Wiley-Interscience
- K. Park and W. Willinger, "Self-similar network traffic: An overview", (2000). Book Chapter, Published
Editor(s): Self-Similar Network Traffic and Performance Evaluation
Bibliography: Wiley-Interscience,
- H. Ren and K. Park, "Toward a theory of differentiated services.", (2000). Proceedings, Accepted
Collection: International Workshop on Quality of Service
Bibliography: IEEE/IFIP
- H. Ren and K. Park, "Efficient shaping of user-specified QoS using aggregate-flow control", (2000). Proceedings, Accepted
Collection: International Workshop on Quality of Future Internet Services
Bibliography: Lecture Notes in Computer Science
- T. Tuan and K. Park, "Multiple time scale redundancy control for QoS-sensitive transport of real-time traffic", (2000). Proceedings, Published
Collection: INFOCOM '00
Bibliography: IEEE

T. Tuan and K. Park, "Congestion Control for self-similar network traffic", (2000). Book Chapter, Published
Collection: Self-Similar Network Traffic and Performance Evaluation
Bibliography: Wiley-Interscience

Husni Fahmi, Walid G. Aref, Muddasar Latif, and Arif Ghafoor, "Distributed Framework for Real-Time Multimedia Object Communication", (2000). Proceedings, Published
Collection: Third International Symposium on Object-Oriented Real-Time Distributed Computing
Bibliography: IEEE

Jianping Fan, Walid G. Aref, and A. Rezgui, "Automatic Object-Oriented Segmentation and Conceptual Modeling Algorithm towards Content-Based Representation", (2000). Proceedings, Accepted
Collection: International Conference on Visual Communication and Image Processing
Bibliography: 2000

Elisa Bertino, Mostafa Hammad, Walid G. Aref, and Ahmed Elmagarmid, "Access Control in Video Database Systems", (2000). Proceedings, Submitted
Bibliography: 2000

W. Aref and I. Ilyas, "The SP-Tree: A generalized space-partitioning Indexing structure", (2000). Proceedings, Submitted
Bibliography: Unknown

W. Aref and William J. McIver, "Indexing Techniques for Video Database Systems", (2000). Survey Article, Submitted
Bibliography: Unknown

S. Dagtas and A. Ghafoor, "Indexing and Retrieval of Video Based Spatial Relation Sequences", (1999). Proceedings, Published
Collection: Multimedia '99
Bibliography: ACM

W. Al-Khatib and A. Ghafoor, "An Approach for Video Data Modeling and Query Processing", (1999). Proceedings, Published
Collection: Multimedia '99
Bibliography: ACM

W. Al-Khatib, "Modeling and Querying in Video Databases", (2001). Thesis, Published
Bibliography: Doctoral Thesis, School of Electrical and Computer Engineering

K. Park and W. Willinger, "The Internet as a Large-Scale Complex Systems.", (). Book, Accepted
Editor(s): K. Park and W. Willinger
Collection: The Internet as a Large-Scale Complex Systems.
Bibliography: To be published by Oxford Press, 2001.

K. Park and H. Lee, "On the effectiveness of probabilistic packet marking for IP traceback under denial of service attack.", (2001). Conference Publication, To appear in Proc. IEEE INFOCOM '01, 2001.
Bibliography: Submitted

H. Ren and K. Park, "Optimal aggregate-flow scheduling.", (2001). Conference Publication, Accepted
Collection: Proc. SPIE International Conference on Scalability and Traffic Control in IP Networks
SPIE International Conference on Scalability and Traffic Control in IP Networks
Bibliography: 2001

H. Ren and K. Park, "Efficient shaping of user-specified QoS using aggregate-flow control.", (2000). Workshop Paper, Published
Collection: Proc. International Workshop on Quality of Future Internet Services

Bibliography: Lecture Notes in
Computer Science, pp. 259-271, 2000.

H. Ren and K. Park., "Toward a theory of differentiated services.", (2000). Workshop Paper, Published
Collection: Proc.
IEEE/IFIP International Workshop on Quality of Service,
Bibliography: pp. 211-220

J. Cruz and K. Park, "Towards communication-sensitive load balancing.", (2001). Conference Paper, Accepted
Collection: Proc. IEEE International Conference on Distributed Computing
Systems
Bibliography: 2001

A. Vakali and C. Stupa, "A QoS based Disk Subsystem", (2001). Conference Paper, Accepted
Collection: Proceedings of the 6th International Conference on Computers and Their Applications (CATA-2001)
Bibliography: 2001

Elisa Bertino, Moustafa A. Hammad, Walid G. Aref and Ahmed K. Elmagarmid., "Access Control Model for Video Database Systems.",
(2000). Conference Paper, Published
Collection: Proceeding of CIKM, Ninth international conference on information and knowledge management
Bibliography: 2000

Mohand-Sa'd Hacid, Abdelmounaam Rezgui, and Jianping Fan., "Geometric Shape Representation for Image Retrieval", (2000). Workshop
Paper, Published
Collection: Proceedings of the Sixth International Workshop on Multimedia Information Systems (MIS'2000)
Bibliography: pages: 156-165.

Jianping Fan, Walid G. Aref and A. Rezgui, "Automatic Object-Oriented Segmentation and Conceptual Modeling Algorithm Towards
Content-Based Representation", (2000). Conference Paper, Published
Collection: International Conference on Visual Communication and Image Processing
Bibliography: 2000

Jianping Fan, Mohand S. Hacid, X. Zhang and Ahmed K. Elmagarmid, "Adaptive Motion-compensated Semantic video object extraction towards
content-based indexing.", (2000). Conference Paper, Published
Collection: International Conference on Internet and Multimedia Systems and Application
Bibliography: pp.430-435, Nov.19-23

Jianping Fan, Mohand S. Hacid, David K.Y. Yau and Ahmed K. Elmagarmid, "Model-based semantic object extraction for content-based video
representation and indexing", (2001). Conference Paper, Published
Collection: Proc. SPIE: Storage and Retrieval for Media Databases
Bibliography: 2001

Jianping Fan and Ahmed K. Elmagarmid, "A semi-automatic semantic algorithm for video object extraction and temporal tracking", (2001).
Conference Paper, Published
Collection: Signal Processing: Image Communication
Bibliography: 2001

Mohand-Said Hacid, Evimaria Terzi and Athena Vakali, "Querying XML with Constraints", (2001). Conference Paper, Accepted
Collection: Special Session on XML Data Management and Applications, Proceedings 2nd International Conference on Internet Computing
Bibliography: 2001

A. Vakali and E. Terzi, "A Two-Level Representation Model for Effective Video Data Storage", (2001). Workshop paper, Published
Collection: MIS'2000, Proceedings of the Sixth International Workshop on Multimedia Information Systems
Bibliography: 2001

- Elisa Bertino, Tiziana Catarci, Ahmed K. Elmagarmid and Mohand-Sa´d Hacid., "A Database Approach to QoS Management in Video Databases", (). Conference Paper, Submitted
Bibliography: Submitted
- Ahmed K. Elmagarmid, Jianping Fan, Mohand-Sa´d Hacid and Farouk Toumani, "Discovering Structural Associations in Video Databases", (). Conference Paper, Submitted
Bibliography: Submitted
- Ahmed K. Elmagarmid and Mohand-Sa´d Hacid, "A Constraint Systems for Ontologies", (). Conference Paper, Submitted
Bibliography: Submitted
- Elisa Bertino, Ahmed K. Elmagarmid and Mohand-Sa´d Hacid, "Ordering and Path Constraints over Semistructured Data", (). Conference Paper, Submitted
Bibliography: Submitted
- Edoardo Ardizzone, Ahmed K. Elmagarmid, Jianping Fan and Mohand-Sa´d Hacid., "Semantic Modeling for Video Browsing Systems", (). Book, Submitted
Bibliography: Submitted
- Ahmed K. Elmagarmid, Mohand-Sa´d Hacid and Evimaria Terzi, "A Framework for Approximate Query Answers in XML", (). Conference Paper, Submitted
Bibliography: Submitted
- R. Chari, and S. Prabhakar., "Prefix Caching and Replication: Techniques for Large Scale Multimedia Document Storage", (). Conference paper, Submitted
Bibliography: Submitted
- R. Sion, A. Elmagarmid, S. Prabhakar, and A. Rezgui, "A Database-Centric Approach to Enabling End-to-End QoS for Multimedia Repositories", (). Workshop Paper, Submitted
Bibliography: Submitted
- Sonia Fahmy, Raj Jain, Bobby Vandalore, and Rohit Goyal, "A Framework for Virtual Channel onto Virtual Path Multiplexing in ATM-ABR", (2000). Proceedings, Published
Collection: Proceedings of IEEE GLOBECOM
Bibliography: volume 3, pp. 1724-1730
- Ahsan Habib, Sonia Fahmy, and Bharat Bhargava, "Design and Evaluation of an Adaptive Traffic Conditioner for Differentiated Services Networks", (2001). Proceedings, Published
Collection: Proceedings of IEEE ICCCN
Bibliography: pp. 90-95
- Chin-ying Wang and Sonia Fahmy, "Dynamics of the pgmcc multicast congestion control protocol", (2002). Proceedings, Accepted
Bibliography: Proceedings of the Grace Hopper Celebration (GHC)
- Minseok Kwon and Sonia Fahmy, "A comparison of load-based and queue-based active queue management algorithms", (2002). Proceedings, Accepted
Bibliography: Proceedings of SPIE ITCOM
- Minseok Kwon and Sonia Fahmy, "Topology-Aware Overlay Networks for Group Communication", (2002). Proceedings, Accepted
Bibliography: Proceedings of ACM NOSSDAV
- Minseok Kwon and Sonia Fahmy, "TCP Increase/Decrease Behavior for Explicit Congestion Notification (ECN)", (2002). Proceedings, Accepted
Bibliography: Proceedings of IEEE ICC Symposium on High-Speed Networks

Ahsan Habib, Bharat Bhargava, and Sonia Fahmy, "A Round Trip Time and Timeout Aware Traffic Conditioner for Differentiated Services Networks", (2002). Proceedings, Accepted

Bibliography: Proceedings of IEEE ICC Symposium on Communications QoS and Reliability

Sonia Fahmy, "Network traffic management", (2002). Book chapter, Accepted

Collection: Encyclopedia of Telecommunications

Bibliography: John Wiley & Sons

Sonia Fahmy, "New TCP Standards and Flavors; and The ns Network Simulator", (2002). Book chapter, Accepted

Collection: High Performance TCP/IP Networking

Bibliography: Prentice Hall

X. Q. Zhu, J. P. Fan, A. K. Elmagarmid, W. A. Aref, "Hierarchical video summarization for medical data", (2002). Proceedings, Published

Bibliography: Proceeding of SPIE Storage and retrieval for media database, volume 4676, pp. 395-406

X. Q. Zhu, J. P. Fan, A. K. Elmagarmid, "Towards facial feature locating and verification for omni-face detection in video/images", (2002).

Proceedings, Accepted

Bibliography: 9th IEEE International Conference on Image Processing (ICIP)

W. Aref, A. C. Catlin, A. K. Elmagarmid, J. Fan, J. Guo, M. Hammad, I. F. Ilyas, M. S. Marzouk, S. Prabhakar, A. Rezgui, E. Terzi, Y. Tu, A. Vakali, X. Q. Zhu, "A distributed database server for continuous media", (2002). Proceedings, Published

Bibliography: ICDS'02 Proceedings of the 18th International Conference on Data Engineering, pp 490-491

I. F. Ilyas, W. Aref, and A. K. Elmagarmid, "Joining ranked inputs in practice", (2002). Proceedings, Accepted

Collection: 28th VLDB Conference

Bibliography: Hong Kong, China

M. Hammad, W. Aref, and A. K. Elmagarmid, "Joining multiple data streams with window constraints", (2002). Technical Report, Submitted

Bibliography: CSD TR #02-010, Computer Science Department, Purdue University

E. N. Houstis, G. Balakrishnan, A. Catlin, N. Dhanjani, S. Lalis, M. Stamatogiannakis, J. R. Rice, and C. Houstis, "Network-based Scientific Computing: The Architecture of Scientific Software", (2001). Book chapter, Published

Editor(s): Ronald F. Boisvert and Ping Tak Peter Tang

Bibliography: Kluwer Academic Publishers, Norwell, MA, pp. 3-28

Ronald F. Boisvert and Elias Houstis, "Computational Science, Mathematics and Software", (2002). Book, Published

Bibliography: Purdue University Press, West Lafayette, Indiana, ISBN 1-55753-250-8

E. Houstis and J.R. Rice, "On the Future of Problem Solving Environments", (2002). Book chapter, Published

Editor(s): R. F. Boisvert and E. N. Houstis

Collection: Computational Science, Mathematics and Software

Bibliography: Purdue University Press

E. N. Houstis, A. Catlin, N. Dhanjani, and J. R. Rice, "WebPDELab: A network based problem solving environment for partial differential equations", (2002). Book chapter, Published

Editor(s): Fox and Dongarra

Collection: Problem Solving Environments

Bibliography: to appear

Y. Lu and B. Bhargava, "Achieving Flexibility and Scalability: A New Architecture for Wireless Network", (2001). Proceedings, Published

Bibliography: Proc. of Intl. Conf. on Internet Computing (IC'2001)

- Y. Lu, B. Bhargava, and M. Hefeeda, "An Architecture for Wireless Networking", (2001). Proceedings, Published
Bibliography: IEEE Workshop on Reliable and Secure Application in Mobile Environment
- A. Habib and B. Bhargava, "Unresponsive Flow Detection and Control using the Differentiated Services Framework", (2001). Proceedings, Published
Bibliography: Proc of 13th IASTED PDCS
- D. Xu, M. Hefeeda, S. Hambrusch, and B. Bhargava, "On Peer-to-Peer Media Streaming", (2002). Proceedings, Accepted
Bibliography: ICDS
- Kihong Park, Heejo Lee, "On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law internets. I", (2001). Proceedings, Published
Bibliography: In Proc. ACM SIGCOMM
- Kihong Park, Heejo Lee, "On the effectiveness of probabilistic packet marking for IP traceback under denial of service attack.", (2001). Proceedings, Published
Bibliography: In Proceedings IEEE INFOCOM
- Kihong Park, Tsunyi Tuan, "Multiple time scale TCP: harnessing proactivity in self-similar burstiness", (2001). Proceedings, Published
Bibliography: In Proc. ITC-17
- Husni Fahmi, "Multimedia Document Scheduling using Proxies in a Distributed Environment", (2002). Thesis, Published
Bibliography: PhD Thesis, School of Electrical and Computer Engineering
- J. B. D. Joshi, K. Li, H. Fahmi, B. Shafiq, and A. Ghafoor, "A Model for Secure Multimedia Document Database System in a Distributed Environment", (2002). Special Issue, Accepted
Collection: IEEE Transactions on Multimedia on Multimedia Databases
Bibliography: Volume 4, Number 2
- W. G. Aref and I. F. Ilyas, "An Extensible Index for Spatial Databases", (2001). Proceedings, Published
Collection: The 13th International Conference on Statistical and Scientific Databases
Bibliography: Virginia
- W. G. Aref and I. F. Ilyas, "A Framework for Supporting the Class of Space Partitioning Trees", (2001). Technical Report, Published
Bibliography: CSD-TR #01-002, Department of Computer Sciences, Purdue University
- W. G. Aref, K. El-Bassyouni, I. Kamel, and M. F. Mokbel, "Scalable QoS-Aware Disk Scheduling", (2002). Proceedings, Published
Collection: International Database Engineering and Applications Symposium (IDEAS)
Bibliography: Alberta, Canada
- W. G. Aref, A. Catlin, A. Elmagarmid, J. Fan, M. Hammad, I. Ilyas, M. Marzouk, and X. Zhu, "Search and Discovery in Digital Video Libraries", (2002). Technical Report, Published
Bibliography: CSD TR #02-005, Computer Sciences Department, Purdue University
- M. Hammad, W. Aref, A. Catlin, M. Elfeky, and A. Elmagarmid, "A Stream Database Server for Sensor Applications", (2002). Technical Report, Published
Bibliography: CSD TR #02-009, Computer Sciences Department, Purdue University
- X. Zhu, J. Fan, W. Aref, and A. Elmagarmid, "ClassMiner: Mining Medical Video Content Structure and Events Toward Efficient Access and Scalable Skimming", (2002). Proceedings, Published
Collection: SIGMOD Workshop on Data Mining and Knowledge Discovery
Bibliography: Madison, WI
- M. Hammad, W. Aref, and A. Elmagarmid, "Search-based Buffer Management Policies for Streaming in Continuous Media Servers", (2002).

Proceedings, Accepted
 Collection: IEEE International Conference on Multimedia and Exposition
 Bibliography: Lausanne, Switzerland

K. Frikken, M. J. Atallah, and Sunil Prabhakar, "Optimal Parallel I/O for Range Queries through Replication", (2002). Proceedings, Accepted
 Collection: 13th International Conference on Database and Expert Systems Applications
 Bibliography: DEXA

D. V. Kalashnikov, Sunil Prabhakar, W. G. Aref, and S. E. Hambrusch, "Efficient Evaluation of Continuous Range Queries on Moving Objects", (2002). Proceedings, Accepted
 Collection: 13th International Conference on Database and Expert Systems Applications
 Bibliography: DEXA

W. G. Aref, A. K. Elmagarmid, J. Fan, J. Guo, M. Hammad, I. Ilyas, M. Marzouk, S. Prabhakar, A. Rezgui, A. Teoh, E. Terzi, Y. Tu, A. Vakali, and X. Zhu, "A Distributed Database Server for Continuous Media", (2002). Proceedings, Accepted
 Collection: 18th International Conference on Data Engineering
 Bibliography: ICDE

Jiangtao Li and Sunil Prabhakar, "Data Placement for Tertiary Storage", (2002). Proceedings, Accepted
 Collection: 10th NASA Goddard Conference on Mass Storage Systems and Technologies/19th IEEE Symposium on Mass Storage Systems
 Bibliography: MSS

S. E. Hambrusch, C.-M. Liu, W. Aref, Sunil Prabhakar, "Query Processing in Broadcasted Spatial Index Trees", (2001). Proceedings, Accepted
 Collection: 7th International Symposium on Spatial and Temporal Databases
 Bibliography: SSTD

Web/Internet Site

URL(s):

<http://www.cs.purdue.edu/msi>

Description:

The MSI site connects all the researchers and associated organizations working on this project. Additional URLs, linked from the MSI web page above, include...

Indiana Telemedicine Incubator

<http://www.cs.purdue.edu/icds/iti>

Indiana Center for Database Systems

<http://www.cs.purdue.edu/icds>

Distributed Multimedia Systems Lab

<http://shay.ecn.purdue.edu/~dmultlab>

Network Systems Lab

<http://www.cs.purdue.edu/nsl>

Center for Education and Research in Information Assurance and Security (CERIAS)

<http://www.cerias.purdue.edu/>

Other Specific Products

Product Type: Data or databases

Product Description:

EduMed database application, a medical education database, was created with the Indiana University School of Medicine to serve their regional Schools of Basic Medical Education.

Sharing Information:

Intellectual Property is currently being established for the application as it is to be distributed to all 9 of Indiana Universities Schools of Basic Medical Education. Once Intellectual Property is established the software will be distributed to the other partners. Subsequent distribution options will be explored and the application is further developed.

Contributions

Contributions within Discipline:

The RTDMDS architecture for the MSI framework 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 PIs for the QoS-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 potential of significantly advancing Web-based multimedia information technology. In addition, the new techniques for multimedia storage management and access control will allow the development of high performance, secure Web-based multimedia servers.

The network QoS architecture contains the some of the most advanced theory developed for differentiated services, with the potential to impact the design and implementation of the Next Generation Internet.

The Q-Bahn QoS testbed, an integrated wired/wireless private testbed comprised of an IP-over-SONET QoS backbone and a WLAN (IEEE 802.11b) based wireless mobile access network, is a one-of-a-kind state-of-the-art research testbed. Its unique features include the Cisco 7206 VXR routers that run a custom IOS (Cisco's router operating system), called purdue-ios, which implements the QoS switching algorithms developed at the Network Systems Lab.

The network security architecture, in addition to its support of standard end-to-end security services, incorporates a new distributed denial of service (DDoS) attack protection component, called route-based distributed packet filtering (DPF), which was advanced in the Network Systems Lab to affect scalable protection against DDoS attacks. It is one of the very few measurement-based (using NLANR AS topology measurements) benchmarked DDoS protection mechanisms in existence, with incremental deployability as one of its key features.

Contributions to Other Disciplines:

Several science and engineering disciplines can benefit from our research related to the MSI framework. In particular, the results and the products from the RTDMDS architecture can be directly extended to the GIS-based applications for the civil engineering discipline, to electronic publication and management of manuals and documents in industrial and mechanical engineering, and to telemedicine and digital library related science and engineering disciplines.

The optimal aggregate-flow switching theory in the networking area significantly advances the state-of-the-art in scheduling (electrical engineering, operations research, industrial engineering, applied mathematics), and the network QoS game theory advances the state-of-the-art in noncooperative game theory (economics, management science, applied mathematics).

The new Q-Bahn testbed is expected to overcome the 'QoS barrier' of the Internet, where few (if any) existing technologies can be expected to be realized on the Internet in the foreseeable future. By taking an intra-domain, incremental approach, the technologies captured by the Q-Bahn testbed can be easily exported to large enterprise and institutional networks where effective QoS support is in significant need.

The DDoS protection architecture based on route-based DPF encompasses several state-of-the-art aspects of scientific research, including fault-tolerance, complexity, and power-law graph connectivity. The route-based DPF results showing effective proactive and reactive DDoS protection is the only known research work, which shows that power-law topology in the Internet autonomous system connectivity context has profound impact on network performance. Collaborative efforts with mathematicians and physicists who have begun studying power-law graphs from other perspectives are being explored as part of this effort.

As an example, Park (co-PI) has given an invited presentation at IPAM (NSF supported center at UCLA) on this very topic to an interdisciplinary audience. Park has also given invited presentations at the Santa Fe Institute on the game theoretic aspects of the QoS part of the project, including a presentation entitled 'The Internet as a Complex System' to SFI Business Network members who are comprised of CTOs and managers of major US and international institutions. Park also organized (with W. Willinger) a joint SFI/NSF workshop entitled 'The Internet as a Large-Scale Complex System' at SFI, March, 2001, which explored the multifaceted nature of the Internet and networked systems, which the MSI project, in a more narrow focus on integrating content/content management (i.e., DB), QoS, and security, is targeted at. We believe this to be just a beginning, albeit an important first step.

Advances in multimedia database technology and high performance networking applied to searchable digital libraries for medical video has the potential to transform (1) basic and continuing medical education, (2) remote diagnostic medicine, and (3) clinical trials management. These three areas of telemedicine have been selected as trial environments for the application of research in multimedia databases through the MSI initiative. Successful demonstrations of multimedia database technology applied to medical education video databases have been presented in late 2001 and early 2002.

Contributions to Human Resource Development:

Both our graduate students (Ph.D. and Master's) and undergraduate students get exposed to cutting edge material, encompassing both theory and implementation, related to the design and development of advanced multimedia database systems, broadband multimedia networking technologies, and information security.

In the area of databases, the students learn the state-of-the-art techniques for storing, accessing, and delivering multimedia documents in a distributed environment. This effort is further facilitated by the Indiana Center for Database Systems under the leadership of Ahmed Elmagarmid (PI).

In the networking area, the leading edge research conducted under the umbrella of the Network Systems Lab, Director Park (co-PI), allows students to understand and experiment with the novel techniques needed for the design and development of the Next Generation Internet and its QoS provisioning architecture. As all components are realized in prototype systems benchmarked over the Q-Bahn testbed, we provide a truly unique and at-the-forefront environment for research and education.

In the security arena, CERIAS, under the leadership of Gene Spafford (co-PI), a well-known center at Purdue University devoted to promoting research and education in information security, has been providing much needed research support and educational opportunities in the this area. One of the features of CERIAS is its interdisciplinary nature, wherein the concerns of areas outside the scope of science and engineering relevant to information security are incorporated.

Contributions to Resources for Research and Education:

Collaborative efforts, in particular, with industrial organizations foster technology transfer as well as collaborative research at the national and international scale. As a part of this project, we have been extremely successful in establishing collaboration with several leading organizations, including Cisco, Sprint, Siemens, and many Indiana-based telemedicine and database companies. We also have close interaction with several academic organizations as listed at the beginning of the report.

Our research and development efforts in multimedia database technology have produced some of the most advanced techniques and models currently available in streaming video database management. This R&D initiative is now the central focus for collaborations with (1) a database company to video-enhance commercial database software and (2) the Indiana University School of Medical to deliver video query and retrieval for medical videos directly into medical school classrooms.

Specifically, Micro Data Base Systems, Inc. (MDBS), a company that markets the TITANIUM database engine, has used MSI-generated technology to integrate video operations into their flagship database product. This has led to further market opportunities for MDBS in the area of video-enhanced database applications. In partnership with the Indiana University School of Medicine, the MSI database technology group is developing a searchable digital medical video library. This library will be accompanied by a database query interface ('a medical education learning environment') that provides online content-based access, query, browse, retrieval, and streaming presentation of indexed medical videos. This MSI-generated technology will deliver real-time online search, discovery, and presentation capabilities for video data directly into medical classrooms.

Contributions Beyond Science and Engineering:

The success of the Next Generation Internet is one of the most pressing and important problems facing technological advancement of society today. The knowledge and experience gained from the network QoS project can have a significant impact on the development of the core international infrastructure critical to all other electronic infrastructures. The integration of secure multimedia document capabilities with this network infrastructure will have significant economic and social impact--and will help protect the environment by reducing our dependence on distribution media for videos, software, and data. At an even broader level, our focused efforts on achieving an integrated environment encompassing databases, networking, and security may provide the foundations and insights into affecting a truly interdisciplinary approach to studying the Internet and its multifaceted features, which is expected to influence all aspects of societal functions.

Special Requirements

Special reporting requirements: None

Change in Objectives or Scope: None

Unobligated funds: less than 20 percent of current funds

Animal, Human Subjects, Biohazards: None

Categories for which nothing is reported:

We continue to make progress in all areas of the MSI project.

Equipment purchases to date, summarized by project area, include:

Networking		
	Cisco routers (9) for QoS experiments	\$234,320
	Intel workstations (24)	\$58,969
	Sun workstations (2)	\$8,986
	Intel-based portables (9)	\$20,556
	Handheld devices (9)	\$6,243
	Networking equipment	\$15,661
Databases and Storage		
	Sun E450/A1000 database server	\$54,950
	Intel-based storage server engine (3)	\$19,188
	Dell-based multimedia proxy servers (4)	\$25,892
	Intel-based workstations (2)	\$19,752
	Intel-based portables (2)	\$6,647
	Sun-based workstations (3)	\$37,372
Image Capture		
	Sony high resolution camcorders	\$7,110
	Futuretel MPEG-2 encoder	\$5,520
Image Compression		
	Intel-based workstation	\$1,466
	SGI high resolution display	\$2,335
Applications		
	Laptop computers (3) for distance learning	\$10,146
	Workstation/printer accessories	\$5,607
	Videoconferencing stations (3) and streamer	\$14,338
Total		\$555,058

We have been extremely successful in obtaining research funds and grants from numerous sources. Noted among them is an award 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.

In addition to our funding from Telcordia, we have also established research collaboration in the area of multimedia document delivery with Siemens Corporation. These efforts have been carried out under the leadership of Ahmed Elmagarmid (PI).

We have also 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. The CERIAS center at Purdue University, which promotes information security research and education, was started by a multi-million dollar grant from the Eli Lilly Foundation. The leadership for this effort is

provided by Gene Spafford, its director, who is also a co-PI. Gene Spafford has also received separate NSF funding for his security research.

Our group has received two large grants, one from NCR/Wal-Mart and the other from Hewlett Packard (HP). The NCR grant is for a 1.6 terabyte parallel database engine valued at approximately \$7 million. This system not only augments the other storage and compute servers we are purchasing through the NSF grant, but also includes a large sample dataset for warehousing and mining experiments. This 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 a National Science Foundation ITR grant (joint with Ohio State University and the University of Illinois, who are subcontractors) for studying multiple time scale traffic control aspects of the network research component, a project led by Park (co-PI).

Park has been awarded a DARPA grant from the ATO FTN program for studying scalable DDoS protection solutions, a component of the security architecture of the MSI project. Park has also received a gift grant from the Xerox Foundation to study the relationship between quality of service, network security, and fault-tolerance.

We received funding from the Santa Fe Institute (SFI) and National Science Foundation to hold 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>.

Sunil Prabhakar (senior personnel) has 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. The co-PIs and personnel include several NSF CAREER awardees and a Presidential Young Investigator awardee.

The remainder of this section highlights specific activities in each of the three major areas of the project: multimedia databases and storage management systems, networking, and security.

Multimedia Databases and Storage Management Systems

There are four major research activities that we are currently pursuing:

1. *Enabling content-based access to video data.* We are pursuing activities aimed at developing a comprehensive video database management system that supports content-based retrieval of large-scale video data as a critical component of multimedia document management system. We are developing techniques and tools for content-based retrieval, semantic browsing, and querying of video, including techniques for (1) partitioning video source into meaningful segments to support more effective video representation and indexing, (2) grouping the physical segments into semantically richer units, (3) detecting

and recognizing faces, and (4) developing metadata to support fast and effective search and browsing techniques for large-scale video. We have incorporated the MPEG-7 representation of multimedia content into our system, which has significantly affected our schema representation and query management. We are developing a query manager that parses and executes MPEG-7 queries. We are also developing an MPEG-7 wrapper to import pre-extracted features in the MPEG-7 format and export features from the database in MPEG-7 document format. To access/query video data and experiment/test system functionality, we have developed a query interface that supports different kinds of queries, such as query by example image (matching image features against extracted video frame features or aggregate shot features) and query by motion (pan, tilt, rotate, zoom, etc.). We are building a web client as a web service to web-enable the entire system.

2. *Adapting traditional databases to handle video data.* The large-scale data volume and continuous media properties of video have a significant impact on the design of video database management systems (VDBMS) in the areas of storage, buffering, query processing, indexing and schema representation. We are developing a prototype video database management system that is capable of handling raw video data as well as its associated metadata. The system uses SHORE, a system developed at the University of Wisconsin, as a storage manager, and PREDATOR, which is a value-added server for SHORE. We continue to address the issue of indexing video feature data to support similarity searching. Our approach is to develop dimension reduction techniques in combination with semantic clustering by exploiting several types of features and the semantics of the video data. We are enhancing the query processing capability of the system to support query by multiple examples and multiple features, and we are developing global similarity ranking techniques to rank results based on any number and combination of features. We have developed mechanisms for executing online queries on multiple video streams for the purposes of video editing, analysis and monitoring. In the query engine we are implementing novel join techniques for online sensor data in the query engine.

3. *Delivering Distributed Multimedia Documents with Support for End-to-End Quality-of-Service.* Currently, we are exploring different approaches that will allow mapping of the user-specified Quality of Presentation (QoP) parameters to Quality of Service (QoS) requirements for different system components of the overall MSI architecture, including storage, servers, networking, and security subsystems. The implementation of the translation mechanisms will be an integral part of the QoS-based resource scheduling modules that will be implemented using several dynamic and static approaches. We have designed an architecture for a real-time distributed multimedia database system (RTDMDS) for managing multimedia documents with Quality of Service (QoS) guarantees. This system is one of the key components of MSI. It allows distributed users to author, store, query, and retrieve multimedia documents over a broadband network. We have developed several modules of RTDMDS and continue the development of the entire system using the equipment acquired through the NSF grant. We are also investigating QoS-sensitive storage of video data.

We are currently developing an end-to-end resource management framework using proxy servers for a large-scale distributed multimedia database system. The system uses a suite of proxies to provide support for interactive function on multimedia presentation. A quality of service (QoS) routing technique has been developed to allow dynamic mapping of QoS connections to network resources, including end-to-end link capacities and buffering capabilities of the proxy servers. The primary objective is to minimize the cost of transmission while simultaneously balancing the load among proxies. Load balancing among proxies is critical because when a proxy becomes heavily loaded as a result of severe load imbalance, the proxy may not be able to meet its QoS commitments to clients. The proposed algorithm is based on a Lagrangian relaxation and a Rounding Breadth First Search algorithm. The method provides a near-optimal solutions in real-time. It is scalable to a large-scale, proxy-based networking infrastructure, and has computation time within the time scale of network state updates of existing network services.

4. *Developing techniques for the large-scale storage of video data.* For physical storage management of multimedia documents, we have designed several novel data placement and scheduling schemes. These schemes are currently being implemented on a Sun E450 server and a Sun A1000 Raid array acquired through this grant. Also, due to the unique temporal nature of video data, we are analyzing several techniques for real-time disk scheduling. Managing large volumes of data necessitates the use of cheap tertiary storage. Given the very high random access cost of tertiary storage, efficient management of data is critical for performance. We are therefore developing data placement, migration, prefetching, caching, and scheduling schemes for the effective retrieval of video from secondary and tertiary storage.

As part of this project, we are currently developing an XML-based multimedia application using the Oracle XML DB which is a high-performance, native XML storage and retrieval technology available with Oracle9i Database Release 2. This development framework complies fully with the W3C XML data model, and absorbs the XML document model into the database for navigation and query. In addition to the native XML repository to the database, the Oracle XML DB encompasses both SQL and XML in a highly interoperable manner. With this combination, the application can store both structured and unstructured data that can be accessed using either SQL or XML operations interchangeably.

Networking

1. 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.
2. We have benchmarked the new architecture and theory using QSim, our ns-based WAN QoS simulator, confirming the theoretical predictions. Our work shows that differentiated services, following our architecture, can provide scalable and efficient user-specified services.

3. We have installed the Purdue Infobahn comprised of four Cisco 7206 VXR routers that form an IP-over-SONET QoS backbone. We have implemented initial signaling and end-to-end controls using both RSVP (Int-Serv) and AS (Diff-Serv) to ascertain the routers' operating capabilities.
4. We have performed LAN-scale testing and benchmarking over the IP-over-SONET IP routers differentiated service provisioning, guaranteed service provisioning, and multiple time scale traffic control. Twelve PCs have dedicated connections to the testbed and are engaged as QoS flows and background traffic generators.
5. The present QoS testbed has been renamed to Q-Bahn (QoS Infobahn) to highlight its broad, comprehensive nature with respect to providing a complete solution for scalable QoS provisioning in IP internets. Q-Bahn consists of 9 Cisco 7206 VXR routers connected as an IP-over-SONET backbone, with 50+ PCs and workstations directly connected to the backbone for benchmarking. The physical topology of Q-Bahn is the same as that of Internet2/Abilene, albeit completely dedicated for advanced QoS and security research, which Internet2/Abilene (due to its operational demands) is not. An even more distinctive feature of Q-Bahn is the operating system running inside the Cisco 7206 IP routers, called *purdue-ios*, which is a custom IOS (Cisco's router operating system) implementing the QoS switching algorithms developed in the Network Systems Lab. This unique collaborative effort between academia and industry led by Park (Director of Network Systems Lab and co-PI of MSI project) allows advanced QoS research to be prototyped in a state-of-the-art production environment, a characteristic strength of the QoS component of the MSI project.
6. The Q-Bahn QoS testbed has been extended to incorporate a wireless mobile component, where a private six access point WLAN cellular network covering the three floors of the CS building is directly connected to the Q-Bahn testbed. The wireless mobile extension has been deemed necessary given the explosion, and expected dominance, of WLAN based local access technology, which injects additional complexities into provisioning effective end-to-end services that must be explicitly incorporated in the global system architecture. The wireless mobile access network is driven by a collection of in-house QoS-enabled applications developed in the Network Systems Lab, one of which, called QVI, a VoIP application incorporating QoS amplification mechanisms, runs over both Windows XP (PCs and laptops) and Windows CE (Compaq iPAQ Pocket PC). In addition to providing QoS-sensitive real-time VoIP and CD quality audio communication over integrated wired/wireless environments, the application QoS amplification infrastructure provides completely transparent QoS support for legacy applications via a QoS module, called Q-Driver, which resides in NDIS of the Windows XP and CE microkernels. Thus, our QoS platform is truly portable to environments characteristic of a typical end user.
7. We have generalized the unified differentiated services framework to a queuing framework (m-class G/G/1 queuing system), which will provide the most comprehensive

theory of aggregate-flow scheduling yet and a significant advancement of scheduling theory.

Security

1. We are focusing our efforts on building 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. At the lower levels, we are building a toolbox for extracting desired features from the underlying video streams. For multimedia document systems, we are developing a security framework that allows the integration of heterogeneous access control policies in a distributed environment.
2. We have advanced and analyzed a denial-of-service (DoS) attack prevention framework based on probabilistic packet marking (PPM), and shown its effectiveness in the presence of single-source DoS attacks. We have used an adversarial framework to derive solutions to optimal decision making in a two-player environment comprised of the attacker and victim/target.
3. We have investigated the distributed DoS (DDoS) attack problem under PPM and general network topologies where the attacker can pick the location and number of attack hosts. We have shown that uncertainty factor amplification --- a measure of the attacker's ability to hide his true location(s) --- is achievable by reducing the attack volume at each individual attack site.
4. Our new work on scalable DDoS attack protection based on route-based distributed packet filtering (DPF) is being recognized as one of the very few effective and deployable DDoS protection mechanisms. One of our future goals is to incorporate the packet filtering mechanisms in the switching elements of the Q-Bahn testbed, thus providing further integration of QoS and network security. Two avenues we seek to explore are extension of the collaborative work with Cisco to include network security and exploration of Intel's IXP 1200 network processor as a programmable IP router platform wherein both QoS and security mechanisms can be implemented, tested, and benchmarked as part of the Q-Bahn testbed.
5. We have investigated the network security architecture of Cisco's router operating system IOS with the aim of incorporating the adaptive security architecture AdSec into its structure.

The findings in this section are organized according to the three major areas of the project: multimedia databases and storage management systems, networking, and security.

Multimedia Databases and Storage Management Systems

1. Content-Based Access to Video

Video Content Analysis. We have developed a novel video shot detection and key frame extraction technique that automatically adapts the threshold for detecting shots based on activities within video sequences. We have also developed a seed region aggregation and temporal tracking procedure to exploit video objects and their trajectories. We have designed and implemented a video preprocessing system with feature extraction functionality. The features currently extracted by our system are color histogram and camera motion. The HSV color space is used to extract the color histogram, and the dimensions of its components are: H (quantified with 18 levels), S (quantified with four levels) and V (quantified with four levels). Nine types of camera motion are extracted and classified as camera motion features. Using this system, the video can be processed with shot segmentation, camera motion classification, hierarchical browsing, and video feature extraction. Based on the results of the video analysis, two categories of video retrieval can be performed with the system: (1) camera motion based video retrieval and (2) key frame based video retrieval. For motion based retrieval, users select a type of camera motion, and the system retrieves shots that fit the specified motion type. For key frame based retrieval, the user first browses the video collection to find an interesting frame to use for submitting an image-matching query. The user submits a query based on the frame and selected frame features. The system extracts the feature of the input frame online, and then uses the distance evaluation function to find all similar shots in the video database.

We have developed algorithms to group physical shots into semantically richer units, such as video groups and video scenes, where each group or scene contains one *story unit*. This process generates a hierarchical video summarization that supports video skimming, browsing, and content access. A user can then browse the scene information and select a story unit to unfold details within the chosen unit.

A facial feature-based omni-face detection/tracking algorithm has been developed. The algorithm uses color segmentation, facial feature filtering, and face candidate verification to support the efficient, effective detection of faces in different poses, orientations and views.

Metadata Presentation. Our system extracts new low-level features from the video that are specified in the MPEG-7 standard, including dominant color, scalable color, homogeneous texture, and edge histogram. Since significant development efforts have been directed towards the specification of the metadata (low-level and high level features) in a format that follows the MPEG-7 description scheme, we have developed a

wrapper to import user-supplied MPEG-7 documents that are generated using multimedia description schemes for high level and low-level feature information. The document features are parsed and mapped to the system relational feature schema. The wrapper also exports feature information from the database as MPEG-7 documents. For the implementation of the wrapper, we use XML-DBMS and Java Packages for document transferral. We are also developing a JDBC driver for PREDATOR for database connectivity.

Browsing Video Data. We have developed an automated technique that combines manual input and knowledge produced by an automatic content characterization technique to build higher-level abstractions of video content. The method is based on description logics for automatically discovering structural associations within and between video sequences. The generated abstraction of the contents of video sequences can be used to navigate the graph structure and determine whether the contents are relevant for the user—thus saving users time by avoiding the unnecessary downloading of large files.

Querying Semi-Structured Data. We designed a constraint-based framework for querying semi-structured data. We provided a class of path constraints of interest in connection with both structured and semi-structured databases. Our constraint language is inspired by Feature Logics. Feature descriptions are used as the main data structure of so-called unification grammars, which is a popular family of declarative formalisms for processing natural language. It provides a partial description of abstract objects by means of functional attributes called features. On top of the constraints we allow the definition of relations by means of definite clauses. The query language is based on the general scheme for handling clauses whose variables are constrained by an underlying constraint theory. Constraints can be seen as quantifier restrictions as they filter out the values that are assigned to the variables of a clause in any of the models of the constraint theory. The language is a hybrid one in the sense that it combines clauses with path constraints. It has a clear declarative and operational semantics.

We have also developed a multimedia authoring, querying, and presentation tool (MAQP) to compose, query, and display multimedia documents. This tool provides a graphical user interface to compose a multimedia document and specify the QoP attributes of its objects along with their temporal ordering in the document.

Web Client. We are working to web-enable the functionality of the entire video database management system. We have built the client as a web service using Microsoft .NET as the platform. We have addressed the missing functionality in current web services platforms and protocols for supporting video and multimedia applications.

2. Managing Video in Databases

Video Database (VDBMS) Prototype. We have introduced the real-time streaming of video as an integrated function within the database system. We are studying issues such as providing fixed bandwidth for each request and modifying buffer management to

handle real-time as well as non real-time requests. We have identified some bottlenecks in current DBMS designs that are not suitable for processing video requests. We have also experimented on the current system under various workloads and conditions, such as large volumes of concurrent requests for different streams. We are investigating the scalability of our system to a growing number of stream requests. Another important aspect is the logging and recovery of video data. This requirement is a consequence of storing the raw video data inside the DBMS. We are investigating current limitations of logging and recovery with huge amounts of storage and long updates to binary contents. We have performed experiments to identify the effect of physical logging of video updates on overall system performance.

High-Dimensional Data. In the VDBMS prototype, we have implemented the GiST indexing framework to realize the R*-tree and the SR-tree as our high-dimensional index structure. In the literature, the SR-tree outperforms other R-tree variants. We are now investigating the development of a general framework for high-dimensional flexible indexing that chooses the most suitable index structure for video data. We are also investigating the use of semantic data clustering (using seed-based clustering algorithm) to build a semantic-aware index structure for video. Semantic clustering can serve as an efficient dimension reduction technique without losing the underlying semantics of video units.

Query Processing and the Join Operator. We extract many high-dimensional features from the video data, on both the frame and shot levels. For each feature, we use a high-dimensional index structure as the access path for similarity queries on this feature. To enable the database management system for querying videos based on multiple features and querying by multiple examples, the similarity results from each index structure need to be joined in a way that gives a global similarity ranking over all specified features. We have enhanced the query processing capability of the system by introducing a new rank-join query operator to handle these queries. The new join operator is a binary, pipelined operator that advances the optimal aggregation algorithms introduced in the literature by transforming an optimal theoretical solution into one that can be implemented in a practical database engine.

Stream Query Processing. We define a stream as an infinite sequence of data items, where items are appended to the sequence over time. Videos can be considered streams of large objects (video frames and audio-related data). We address the issue of executing online queries on multiple video streams for the purposes of video stream editing, analysis, and monitoring. As an example, we might need to hide the faces of patients in a medical education video of patient consultations in order to protect patient privacy. We address the online processing of video since we may not want to store the preprocessed video, and we do not want to create a special edited version of the video for each patient. As a second example, consider the tracking of persons between two video cameras. Online video preprocessing is used to extract “face” objects that appear in each video stream and determine their identity using face detection/recognition techniques. The result of video preprocessing is two streams of object identities. One can execute a join-query on both streams to retrieve the persons (with the same identity) that appear in both

streams. Due to the possibly unlimited size of the stream, the join operation can be restricted by a window of time during which the join can be performed between the corresponding streams.

The source of streaming data is a sensor that can supply data either continuously or asynchronously. The video stream is represented as sensor data. In the current implementation, the sensor is represented as a user-defined data type with specific interfaces. Specifically, the sensor data type implements three main functions: `InitStream`, `ReadStream`, and `CloseStream`. We have developed a new component, the stream manager, to coordinate the retrieval of streams from the sensors and to supply the corresponding query execution. The stream manager uses the three main stream functions to collect data from the sensor, and user-defined functions can be implemented as part of the sensor data type. Multiple sensors can also be defined using this approach.

In the query processor, we have introduced the `StreamScan` physical operator. The main functionality of this operator is to provide the interface between the stream manager and the query execution plan. Initially, the `StreamScan` contacts the Sensor Table and registers stream requests to the stream manager. This triggers the stream manager to start retrieving sensor values. Based on a request from other query operators, `StreamScan` retrieves a tuple-at-a-time from the buffers of the stream manager for processing by the query execution plan. We have already implemented two novel join techniques for online sensor data in the query engine: the nested-loop window-join and hash window-join. We are currently investigating the issues of (1) query execution with limited amounts of memory, (2) intermediate storage, and (3) scalability (queries with large number of streams). We believe that streams of large objects have additional complexity that imposes further challenges on stream query processing.

Online Video Processing. To address the online processing of video, we have implemented video processing functions such as blurring and fast forwarding. Those functions are executed within the query plan and work on the video in the compressed domain so that it is not necessary to decompress, process, and then re-compress the video to produce the final online results.

3. Quality of Service in Video Databases

Quality-of-Service Aware Repository (QuaSAR) Architecture. We have designed a system architecture that supports user quality-sensitive queries within a database framework. The proposed architecture relies upon the notion of QoS aware interfaces to the various components of the system, such as the network layer and the operating system (encompassing CPU, main memory, and disk storage). These interfaces enable real-time determination of the status of the components with respect to the satisfaction of QoS constraints. In addition, these interfaces will support reservation of resources to guarantee the ability to satisfy the user's requested level of quality. A key component of QuaSAR is the enhanced query processing capabilities in contrast to traditional databases. Based on the content component of the query and the content metadata,

alternative plans are generated for the retrieval of the relevant objects. Each plan is annotated with QoS parameters relevant to each component based upon translation of the user's quality parameters for the given plan. Each of the constraints represented by the annotations are tested through the interfaces and, if necessary, reserve resources. If no feasible plan is found, a negotiation step is invoked to adjust the constraints and re-evaluate the feasibility of the plan.

QoS Constraint Specification Language. A language, in the style of constraint database languages for formal specification of QoS constraints has been developed. We have shown that constrained rules are a valuable tool for specifying and managing quality of service in time-based media databases. We also described an algorithm for deciding whether a query 'matches' a specified quality of service. We show that the satisfaction by the system of the user quality requirements can be viewed as a constraint satisfaction problem, and that the re-negotiation can be viewed as constraint relaxation.

Distributed Proxies. We are developing a distributed proxy architecture to allow QoS-based synchronization and caching of multimedia documents in a network environment. The key feature of this architecture is a centralized load balancing mechanism for the network proxy servers that uses a stochastic scheduling approach. Currently, we are simulating our approach using network traffic traces from the Lawrence Berkeley National Laboratory and the Technical University of Berlin. In addition, we have implemented a media stream synchronization technique for distributed multimedia document servers.

Security Constraints. In the area of multimedia database security, we have proposed a colored Petri-net based formalism to provide a multilevel security mechanism. In particular, the model allows specification of security attributes of multiple levels to manage access control to multimedia documents in a distributed multi-domain environment. The MAQP tool can provide such specifications as a part of the overall QoS requirements.

4. Large Scale Video Storage

We have developed a novel hot prefix caching scheme for continuous media placement across the secondary-tertiary boundary. The key idea is to reserve a portion of secondary storage for storing the initial segments of continuous media objects in lieu of the traditional use as a cache for tertiary storage. These segments serve the purpose of masking the extremely high latency of random access to tertiary storage. In order to reduce jitter during playback of documents that are stored on tertiary storage, we propose the use of full replication. The proposed schemes are tested using a simulation of the system under conditions of concurrent access. Our results show that these two techniques result in significant reductions in the startup latency as well as jitter during playback. We are also investigating placement schemes for tertiary storage based on access patterns that show relationships between documents or objects.

Popularity-based models have been proposed such that multimedia (video) data representation guides data placement on a tertiary storage subsystem. A two-level representation model is considered to capture the frequencies of accesses at external (video objects) and internal (video clips) levels. Video data placement is employed on a tertiary storage topology under three well known placement policies governed by the Organ-pipe, the Camel, and the Simulated Annealing algorithms. The latter approach proves to be the most beneficial for the overall multimedia system performance. Furthermore, QoS has been proposed in storage subsystem management towards effective disk space utilization and request servicing. A QoS based storage model for effective user negotiation in terms of scheduling, redundancy, and number of storage devices has been developed. Users can create their own profile with respect to certain QoS attributes in order to specify their requirements. A hierarchical storage model with data elevation among various levels of the storage hierarchy has been simulated. Algorithms of placement among different levels of storage hierarchy and elevation issues have been investigated.

Networking

Advance foundations of effective differentiated services architecture. We have developed a framework for reasoning about the 'goodness' of various differentiated services architectures, which considers both the per-hop behavior (PHB) and edge control parts when provisioning user-specified QoS-sensitive network services. In the PHB part, we have formulated the problem of optimal aggregate-flow per-hop behavior and solved it by showing the optimal aggregate-flow classifier that dominates all other per-hop behaviors with respect to both efficiency and fairness. We have shown that the optimal aggregate-flow classifier—implementable over IPv4 and IPv6 in modern routers and performing stateless per-hop control scheduling in linear time—satisfies certain properties (called (A1), (A2), and (B)), which allows desired end-to-end QoS to be achieved using edge control, both open-loop and closed-loop. Our framework, analysis machinery, and tools answer pressing questions such as 'what is the loss of power due to flow aggregation?', 'what is the impact of a discrete, bounded label set (DSCP in the DS field) on QoS?', 'what edge control can achieve setting of end-to-end TOS label values to achieve system optimal QoS?' to be answered both qualitatively and quantitatively. 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.

Configure Purdue Infobahn QoS Testbed and Perform Initial Tests and Exploratory Implementations. As a starting point to implementing a full-fledged version of our optimal aggregate-flow classifier inside Cisco 7206 VXR routers with system support from Cisco, 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. The latter include saturation effects stemming from processor scheduling and software switching overhead at the router. 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 properties when incorporating all software and processing overheads associated with routers and end stations.

Advance End-to-end Traffic Control for QoS Shaping. We have developed and implemented end-to-end QoS amplification techniques using adaptive packet-level forward error correction and multiple time scale traffic control. Adaptive FEC (AFEC) allows invariant end-to-end QoS to be exported over variable network conditions stemming from shared bottleneck routers and congestion effects. Since differentiated services—even in the presence of optimal aggregate-flow PHB—is subject to occasional variability due to imperfect QoS protection resulting from abstaining from per-flow reservation and admission control, AFEC can provide QoS amplification over varying degrees of imperfect QoS channels which can further improve the operating range and usefulness of differentiated services on top of scalability. 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. An important consequence of MTS traffic control is the ability to mitigate the cost of feedback traffic controls in high delay-bandwidth product networks that is especially pronounced in high-bandwidth wide area networks. By exploiting predictability structure resident at time scales exceeding the round-trip time (RTT) or feedback loop by an order of magnitude or more, timeliness of reactive can be achieved by bridging the 'uncertainty gap' inherent in reactive actions. Thus, 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. In the latter, the sender interfaces with a real-time MPEG compression board (Optibase I for MPEG I video/audio and Futuretel for MPEG II video/audio), which is fed by high-end digital and analog video cameras.

We have implemented and tested both differentiated services and guaranteed services functionalities on Purdue Infobahn. We have observed that RSVP-based per-flow resource reservation shields user flows from the detrimental effects of cross-traffic whereas for differentiated services the achieved service level is more variable. The influence of processing speed of the router can be significant when subject to high load, and it is a characteristic that needs to be explicitly incorporated when computing end-to-end QoS assurances. The first service parallels the standardization effort under the auspices of the Internet2 QoS Working Group (co-PI Kihong Park is a member of the

Architecture Design Team) which seeks to achieve robust Premium Service using IETF's Expedited Forwarding (EF) specification, however, using conventional per-flow reservation.

We have generalized the unified differentiated services framework to a queueing framework -- m-class G/G/1 queueing system -- which is the most general framework upon which to establish a theory of aggregate-flow scheduling. We have shown that aggregate-flow scheduling theory can be derived from the foundations of per-flow scheduling theory -- the same approach as followed in the previous optimal per-hop behavior work without queueing -- which uses the conservation law (a dot product functional which applies to all work-conserving schedulers and their queueing systems) as the basis upon which further refinements are carried out that incorporate ever more structure induced from the stochastic nature of the arrival process (e.g., m-class M/G/1, G/M/1, and M/G/infinity input). In terms of relevance, the new theory plays a similar role as RMS and EDF do for real-time scheduling where they are shown to be optimal for fixed and variable priority schedulers, respectively.

With the wide-spread deployments of WLANs (IEEE 802.11b) and their expected dominance in the future as a de factor local access technology, it has become imperative to explicitly consider the influence of wireless access and mobility on end-to-end performance. Toward this end, we have extended our research platform to include cellular WLAN access networks, and developed an in-house QoS-enabled Voice-over-IP application, called, QVI, which incorporates two QoS amplification mechanisms: (a) packet-level FEC and (b) TOS field label control. Feature (a) incorporates application transparent packet-level FEC control through implementation of a QoS module (called the Q-Driver) inside NDIS, which allows other legacy applications (not just QVI) to invoke packet-level FEC for QoS protection. Feature (b), also implemented by the Q-Driver, allows the TOS field of legacy applications to be imbued with end-to-end QoS control, thus allowing high priority application streams (i.e., their IP packets) to receive prioritized treatment at cognizant routers (our Cisco routers in the Q-Bahn QoS testbed). A future research item that we plan to attack is the inclusion of the MAC layer, wherein integrated QoS management crossing the various layers, including the MAC layer, can be effectively facilitated. Our QVI applications are implemented, tested and benchmarked in real WLAN environments (a wireless extension of the Q-Bahn testbed, which includes a private WLAN comprised of 6 Enterasys RoamAbout R2 802.11b access points, and a number of Compaq iPAQ Pocket PC H3850 handhelds running our QoS environment over/inside Windows XP and Windows CE.

Security

For video databases, we have developed an access control mechanism that can allow different users to view different contents based on their qualifications and credentials. The units of access control can either be a sequence of one or more frames, or parts of a frame, e.g., objects in a frame. For multimedia document system, we have developed a Petri-net based model for role-based access control mechanism that allows selective viewing of different parts of document by authorized users.

We have advanced a framework for probabilistic packet marking for distributed denial-of-service attack (DoS) prevention that 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 DoS, both single-source and distributed. We use a 2-player adversarial framework to analyze the minimax and maximin optimal strategies of the attacker and victim, and show that PPM is effective at preventing DoS attacks. We show using numerical results the outcome of optimal strategies when using typical Internet based configurations including hop count, topology, spoofing, and marking probability.

We have proposed a new proactive and scalable network architecture for distributed denial-of-service (DoS) attack prevention on the Internet called route-based distributed packet filtering (DPF). We have shown that DPF achieves proactiveness and scalability, and we have shown that there is an intimate relationship between the effectiveness of DPF at mitigating distributed DoS (DDoS) attack and the recently discovered power-law Internet topology. The salient features of this work are two-fold. First, we have shown that DPF is able to proactively filter out a significant fraction of spoofed packet flows and prevent attack packets from reaching their targets in the first place. The IP flows that cannot be proactively curtailed are extremely sparse such that their origin can be localized---i.e., IP traceback---to within a small, constant number of candidate sites. We show that the two proactive and reactive performance effects can be achieved by implementing route-based filtering on less than 20\% of Internet autonomous system (AS) sites. Second, we have shown that the two complementary performance measures are dependent on the properties of the underlying AS graph topology. In particular, we have shown that the power-law structure of Internet AS topology leads to connectivity properties which are crucial in facilitating the observed performance effects. As a DDoS prevention architecture, DPF is able to emulate the IP traceback prowess of probabilistic packet marking, while alleviating the latter's three principal weaknesses: (i) need to inscribe link information in the IP packet header, (ii) reactivity---traceback occurs after the impact of DoS attack has been felt---and (iii) scalability where the effort needed to achieve IP traceback grows proportionally with the number of attack hosts engaged in a DDoS attack. We are exploring incorporation of these security mechanisms inside IOS by extending the scope of our collaboration with Cisco Systems to include network security, and implementation of both the QoS and network security mechanisms on Intel IXP 1200 network processor based platforms, where programmable IP router functionalities can be tested and benchmarked as an extension of Q-Bahn.