

**Final Report for Period:** 09/1999 - 08/2005  
**Principal Investigator:** Elmagarmid, Ahmed K.  
**Organization:** Purdue Research Foundation

**Submitted on:** 01/31/2006

**Award ID:** 9972883

**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 QoP/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.

**Name:** Ghanem, Thanaa

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

**Name:** Elfeky, Mohamed

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

Mohamed Elfeky is involved in incremental, online and merge mining of partial periodic patterns in time series databases. He is also working on record linkage and quality assurance of databases.

**Name:** Tu, Yi-Cheng

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

Yi-Cheng is involved in applying tertiary storage mechanisms to augment the current storage manager in the VDBMS. He has also worked on QoS support for multimedia databases.

**Name:** Ali, Mohamad

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

Mohamed works on the STEAM engine variant of the VDBMS video database management system. The STEAM server supports streaming operations and operates as the stream control interface between arrays of distributed data stream sources and end-user clients that access and analyze the streams.

## 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.

**Name:** Brahmhatt, Kintan

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

Kintan Brahmhatt is working on an extension of the query language for the web-based medical query and retrieval application (which is based on multimedia technology developed through MSI). The extension involves application access to UMLS, a standardized medical ontology.

**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:**

**Name:** Hammuda, Ayman

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

Ayman Hammuda is a consultant involved in writing a web-based multimedia medical query and retrieval application based on the distributed multimedia database infrastructure developed through MSI

**Name:** Arunkumar, Sarosh

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

Sarosh Arunkumar is a computer programmer involved in the extension of the VDBMS video database management system as a testbed facility for video database benchmarking. This will allow the research community to use VDBMS to integrate, test and compare new algorithms for video processing and new database components for video database operations.

**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**

Technology transfer for video-enhanced database capabilities between researchers in the MSI project and programmers for the TITANIUM database engine marketed by 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)  
 Gordon Coppoc (Indiana University School of Medicine)  
 Ronald Steuterman (Dir. of Tech Transfer Institute, Purdue Univ.)

**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:**

MSI project funds supported a graduate course at Purdue University in Distributed Systems, including the installation and use of several distributed computing environments providing advanced control of quality of service in distributed environments. This involved both commercial products (e.g., DCE) and research environments (e.g., D'Agents).

The 490S Topics in Computer Science for Undergraduates course (Spring 2002) involved 8 Purdue University Undergraduates in hands-on research, implementation and experimentation for work related to MSI.

The CS590A Advanced Topics in Database Systems course (Fall 2001) covered new trends in database systems, including multimedia database server design issues, advanced query processing, and other topics related to ongoing research for MSI.

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 engaged in several outreach activities related to the MSI research and development efforts:

- Walid Aref was the program committee chair for the International Symposium on Geographic Information Systems in Atlanta, Georgia. 11-02-2002.

- Ahmed Elmagarmid participated in NSF workshops in Intelligent Data Management and in Digital Government (Chicago, IL and Los Angeles, CA)

- Lectures in Database Technologies were presented to the Purdue University School of Science Honors program for undergraduates. (2002).

- We 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 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 gave 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 gave a tutorial entitled 'Internet Traffic Modeling and its Implications to Network Performance and Control' at Performance 2002, in Rome, Italy. It exposed 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. 328, vol. 8, (2002). Published

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

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. 521, vol. 9, (2000). Published

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. 62, vol. , (2001). Published

David K. Y. Yau and Bharat Bhargava, "Heterogeneous CPU Services using Differentiated Admission Control", *Kluwer Multimedia Tools and Applications Journal*, p. 213, vol. 17, (2002). Published

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. 121, vol. 11, (2002). Published

- 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. 79, vol. 16, (2001). Published
- 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. 4, vol. , (2002). Published
- 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. 35, vol. 30, (2001). Published
- Mohand-Sa'd Hacid, Farouk Toumani and Ahmed K. Elmagarmid., "Constraint-Based Approach to Semistructured Data.", Fundamenta Informaticae Journal, p. 53, vol. 47, (2001). Published
- 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. 521, vol. 9, (2000). Published
- Jianping Fan, Xiangquan Zhu, Mohand-Sa'id Hacid and Ahmed K. Elmagarmid, "Model-Based Video Classification Toward Multi-level Representation, Indexing and Accessing.", Journal of Multimedia Tools and Applications, p. 97, vol. 17, (2001). Published
- 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). Published
- A. Vakali and E. Terzi, "Video Data Storage Policies : An Access Frequency Based Approach", Computers & Electrical Engineering Journal, Elsevier, p. 447, vol. 28, (2001). Published

- 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. 45, vol. 23, (2001). Published
- 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. 150, vol. 78, (2002). Published
- 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. 145, vol. 17, (2002). Published
- J. Fan and A. K. Elmagarmid, "Semi-automatic algorithm for video object extraction and temporal tracking", *Signal Processing: Image Communication*, p. ??, vol. 17, (2002). Published
- 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). Published
- 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. 332, vol. 16, (2004). Published
- X. Zhu, A. K. Elmagarmid, X. Xue, L. Wu, A. C. Catlin, "InsightVideo: Toward hierarchical content organization for efficient browsing, summarization and retrieval", *IEEE Transactions on Multimedia*, p. 648, vol. 7, (2005). Published
- 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. 83, vol. 149, (2001). Published
- 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. 1188, vol. 15, (2003). Published

- 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, 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
- Christopher Clifton, "Change Detection in Overhead Imagery using Neural Networks", *Int. Journal of Applied Intelligence, Kluwer Academic Publishers.*, p. 45, vol. 18(2), (2003). Published
- Sunil Prabhakar, Rahul Chari, "Minimizing Latency and Jitter for Large Scale Multimedia Repositories through Prefix Caching", *International Journal on Image and Graphics (IJIG)*, p. 95, vol. 3(1), (2003). Published
- Radu Sion, M. J. Atallah, Sunil Prabhakar, "On Watermarking Numeric Sets", *International Workshop on Digital Watermarking (IWDW) 2002*, p. 130, vol. , (2002). Published
- Sunil Prabhakar, Y. Xia, D. Kalashnikov, W. Aref, S. E. Hambrusch, "Query Indexing and Velocity Constrained Indexing: Scalable Techniques for Continuous Queries on Moving Objects", *IEEE Transactions on Computers*, p. 1124, vol. 51(10), (2002). Published
- K. Frikken, M. J. Atallah, Sunil Prabhakar, R. Safavi-Naini, "Optimal Parallel I/O for Range Queries through Replication", In *Proceedings of the 13th International Conference on Database and Expert Systems Applications - DEXA 2002, Aix-en-Provence, France*, p. 669, vol. , (2002). Published
- Dmitri Kalashnikov, Sunil Prabhakar, S. E. Hambrusch, W. Aref, "Efficient Evaluation of Continuous Range Queries on Moving Objects", *Technical Report number 02-015, Purdue University*, p. 1, vol. 15, (2002). Published
- Mohamed F. Mokbel, Walid G. Aref, and Ibrahim Kamel, "Analysis of Multi-Dimensional Space-Filling Curves", *GeoInformatica*, p. 179, vol. 7, (2003). Published
- Moustafa A. Hammad, Michael J. Franklin, Walid G. Aref and Ahmed K. Elmagarmid, "Scheduling for shared window joins over data streams", *Technical Report CSD TR#03-001, Purdue University*. Jan. 2003, p. 1, vol. , (2003). Published
- Moustafa A. Hammad, Walid G. Aref and Ahmed K. Elmagarmid, "Joining Multiple Data Streams with Window Constraints", *Technical Report CSD TR#02-010, Purdue University*. May 2002, p. 1, vol. , (2002). Published
- M. Elfeky, W. Aref, M. Atallah and A. Elmagarmid, "Periodicity Detection in Time Series Databases", *Technical Report, Purdue University*, p. 1, vol. , (2002). Published
- B. Armbruster, C. Smith and K. Park., "A packet filter placement problem with application to defense against spoofed denial of service attacks", *European Journal of Operational Research*, p. , vol. , (2005). Accepted
- S. Choi, K. Park and C. Kim, "Performance impact of inter-layer dependence in infrastructure WLANs", *IEEE Transactions on Mobile Computing*, p. 2005, vol. , ( ). Accepted

- A. R. bashandy, E. K. P. Chong, and A. Ghafoor, "Generalized Quality-of-Service Routing with Resource Allocation", IEEE Journal on Selected Areas in Communications, p. 450, vol. 23, (2005). Published
- S. Choi, K. Park and C. Kim, "On the performance characteristics of WLANs: revisited", Proc. ACM SIGMETRICS '05, p. 97, vol. , (2005). Published
- W. A. Sheikh, A. R. Bashandy, and A. Ghafoor, "A Combined Approach for QoS based multicast Routing and Resource Allocation", Proc. IEEE International Conference on Communications, p. 132, vol. , (2005). Published
- Ihab F. Ilyas, Walid G. Aref, "Rank-Aware Query Processing and Optimization", ICDE 2005, p. 1144, vol. , (2005). Published
- Xiaopeng Xiong, Mohamed F. Mokbel, Walid G. Aref, "SEA-CNN: Scalable Processing of Continuous K-Nearest Neighbor Queries in Spatio-temporal Databases", ICDE 2005, p. 643, vol. , (2005). Published
- Mohamed G. Elfeky, Walid G. Aref, Ahmed K. Elmagarmid, "WARP: Time Warping for Periodicity Detection", ICDM 2005, p. 138, vol. , (2005). Published
- Mohamed H. Ali, Walid G. Aref, Cristina Nita-Rotaru, "SPASS: scalable and energy-efficient data acquisition in sensor databases", MobiDE 2005, p. 81, vol. , (2005). Published
- Mohamed F. Mokbel, Walid G. Aref, "GPAC: generic and progressive processing of mobile queries over mobile data", Mobile Data Management 2005, p. 155, vol. , (2005). Published
- Moustafa A. Hammad, Walid G. Aref, Ahmed K. Elmagarmid, "Optimizing In-Order Execution of Continuous Queries over Streamed Sensor Data", SSDBM 2005, p. 143, vol. , (2005). Published
- Mohamed H. Ali, Mohamed F. Mokbel, Walid G. Aref, Ibrahim Kamel, "Detection and Tracking of Discrete Phenomena in Sensor-Network Databases", SSDBM 2005, p. 163, vol. , (2005). Published
- Hicham G. Elmongui, Mohamed F. Mokbel, Walid G. Aref, "Spatio-temporal Histograms", SSTD 2005, p. 19, vol. , (2005). Published
- Mohamed H. Ali, Walid G. Aref, Raja Bose, Ahmed K. Elmagarmid, Abdelsalam Helal, Ibrahim Kamel, Mohamed F. Mokbel, "NILE-PDT: A Phenomenon Detection and Tracking Framework for Data Stream Management Systems", VLDB 2005, p. 1295, vol. , (2005). Published
- Mohamed F. Mokbel, Xiaopeng Xiong, Moustafa A. Hammad, Walid G. Aref, "Continuous Query Processing of Spatio-Temporal Data Streams in PLACE", GeoInformatica, p. 343, vol. 9, (2005). Published
- Mohamed F. Mokbel, Walid G. Aref: PLACE, "A Scalable Location-aware Database Server for Spatio-temporal Data Streams", IEEE Data Eng. Bull., p. 3, vol. 9, (2005). Published
- Mohamed G. Elfeky, Walid G. Aref, Ahmed K. Elmagarmid, "Periodicity Detection in Time Series Databases", IEEE Trans. Knowl. Data Eng., p. 875, vol. 17, (2005). Published
- Mohamed G. Elfeky, Walid G. Aref, Ahmed K. Elmagarmid, "Using Convolution to Mine Obscure Periodic Patterns in One Pass", EDBT 2004, p. 605, vol. , (2005). Published
- Mohamed F. Mokbel, Ming Lu, Walid G. Aref, "Hash-Merge Join: A Non-blocking Join Algorithm for Producing Fast and Early Join Results", ICDE 2004, p. 251, vol. , (2004). Published
- Thanaa M. Ghanem, Rahul Shah, Mohamed F. Mokbel, Walid G. Aref, Jeffrey Scott Vitter, "Bulk Operations for Space-Partitioning Trees", ICDE 2004, p. 29, vol. , (2004). Published

- Mohamed F. Mokbel, Walid G. Aref, Khaled M. Elbassioni, Ibrahim Kamel, "Scalable Multimedia Disk Scheduling", ICDE 2004, p. 498, vol. , (2004). Published
- Moustafa A. Hammad, Mohamed F. Mokbel, Mohamed H. Ali, Walid G. Aref, Ann Christine Catlin, Ahmed K. Elmagarmid, M. Eltabakh, Mohamed G. Elfeky, Thanaa M. Ghanem, R. Gwadera, Ihab F. Ilyas, Mirette S. Marzouk, Xiaopeng Xiong, "Nile: A Query Processing Engine for Data Streams", ICDE 2004, p. 851, vol. , (2004). Published
- Ihab F. Ilyas, Rahul Shah, Walid G. Aref, Jeffrey Scott Vitter, Ahmed K. Elmagarmid, "Rank-aware Query Optimization", SIGMOD Conference 2004, p. 203, vol. , (2004). Published
- Mohamed F. Mokbel, Xiaopeng Xiong, Walid G. Aref, "SINA: Scalable Incremental Processing of Continuous Queries in Spatio-temporal Databases", SIGMOD Conference 2004, p. 623, vol. , (2004). Published
- Xiaopeng Xiong, Mohamed F. Mokbel, Walid G. Aref, Susanne E. Hambrusch, Sunil Prabhakar, "Scalable Spatio-temporal Continuous Query Processing for Location-aware Services", SSDBM 2004, p. 317, vol. , (2004). Published
- Mohamed F. Mokbel, Xiaopeng Xiong, Moustafa A. Hammad, Walid G. Aref, "Continuous Query Processing of Spatio-temporal Data Streams in PLACE", STDBM 2004, p. 57, vol. , (2004). Published
- Mohamed F. Mokbel, Xiaopeng Xiong, Walid G. Aref, Susanne E. Hambrusch, Sunil Prabhakar, Moustafa A. Hammad, "PLACE: A Query Processor for Handling Real-time Spatio-temporal Data Streams", VLDB 2004, p. 1377, vol. , (2004). Published
- Thanaa M. Ghanem, Walid G. Aref, "Databases Deepen the Web", IEEE Computer, p. 116, vol. 37, (2004). Published
- Xingquan Zhu, Xindong Wu, Jianping Fan, Ahmed K. Elmagarmid, Walid G. Aref, "Exploring video content structure for hierarchical summarization", Multimedia Syst, p. 98, vol. 10, (2004). Published
- Walid G. Aref, Ann Christine Catlin, Ahmed K. Elmagarmid, Jianping Fan, Moustafa A. Hammad, Ihab F. Ilyas, Mirette S. Marzouk, Sunil Prabhakar, Yi-Cheng Tu, Xingquan Zhu, "VDBMS: A testbed facility for research in video database benchmarking", Multimedia Systems, p. 575, vol. 9, (2004). Published
- Ihab F. Ilyas, Walid G. Aref, Ahmed K. Elmagarmid, "Supporting top-k join queries in relational databases", VLDB J., p. 207, vol. 13, (2004). Published
- Mohamed F. Mokbel, Walid G. Aref, Susanne E. Hambrusch, Sunil Prabhakar, "Towards scalable location-aware services: requirements and research issues", GIS 2003, p. 110, vol. , (2003). Published
- Xingquan Zhu, Walid G. Aref, Jianping Fan, Ann Christine Catlin, Ahmed K. Elmagarmid, "Medical Video Mining for Efficient Database Indexing, Management and Access", ICDE 2003, p. 569, vol. , (2003). Published
- Mohamed F. Mokbel, Walid G. Aref, Ananth Grama, "Spectral LPM: An Optimal Locality-Preserving Mapping using the Spectral (not Fractal) Order", ICDE 2003, p. 699, vol. , (2003). Published
- Walid G. Aref, Moustafa A. Hammad, Ann Christine Catlin, Ihab F. Ilyas, Thanaa M. Ghanem, Ahmed K. Elmagarmid, Mirette S. Marzouk, "Video query processing in the VDBMS testbed for video database research", MMDB 2003, p. 25, vol. , (2003). Published
- Moustafa A. Hammad, Walid G. Aref, Ahmed K. Elmagarmid, "Stream Window Join: Tracking Moving Objects in Sensor-Network Databases", SSDBM 2003, p. 75, vol. , (2003). Published
- Mohamed F. Mokbel, Walid G. Aref, "On Query Processing and Optimality Using Spectral Locality-Preserving Mappings", SSTD 2003, p. 102, vol. , (2003). Published
- Moustafa A. Hammad, Michael J. Franklin, Walid G. Aref, Ahmed K. Elmagarmid, "Scheduling for shared window joins over data streams", VLDB 2003, p. 297, vol. , (2003). Published

- Ihab F. Ilyas, Walid G. Aref, Ahmed K. Elmagarmid, "Supporting Top-k Join Queries in Relational Databases", VLDB 2003, p. 754, vol. , (2003). Published
- Mohamed F. Mokbel, Thanaa M. Ghanem, Walid G. Aref, "Spatio-Temporal Access Methods", IEEE Data Eng. Bull., p. 40, vol. 26, (2003). Published
- R. Cheng, D. V. Kalashnikov, S. Prabhakar, "Evaluation of Probabilistic Queries Over Imprecise Data in Constantly-Evolving Environments", Information Systems: An International Journal, p. , vol. , (2005). Accepted
- S. Hambrusch, C.-M. Liu, S. Prabhakar, "Broadcasting and Querying Multi-Dimensional Index Trees in a Multi-Channel Environment", Information Systems: An International Journal, p. , vol. , (2005). Accepted
- D. V. Kalashnikov, S. Prabhakar, "Fast similarity join for multi-dimensional data", Information Systems: An International Journal, p. , vol. , (2005). Accepted
- Lou, K., Iyer, N., Jayanti, S., Kalyanaraman, Y., Sunil, P., and Ramani, K, "Supporting Effective and Efficient Three-Dimensional Shape Retrieval", Journal of Engineering Design, p. 175, vol. 16, (2005). Published
- Yi-Cheng Tu, Jianzhong Sun, Mohamed Hefeeda, S. Prabhakar, "An Analytical Study of Peer-to-Peer Media Streaming Systems", ACM Transactions on Multimedia Computing, Communications, and Applications (ACM TOMCCAP), p. , vol. , (2005). Accepted
- R. Sion, M. J. Atallah, S. Prabhakar, "Rights Protection for Categorical Data", IEEE Transactions on Knowledge and Data Engineering (TKDE), p. 912, vol. 17, (2005). Published
- R. Sion, M. J. Atallah, S. Prabhakar, "Rights Protection for Relational Data through Watermarking", IEEE Transactions on Knowledge and Data Engineering (TKDE), p. 1509, vol. 16, (2004). Published
- R. Cheng, D. V. Kalashnikov, S. Prabhakar, "Querying Imprecise Data in Moving Object Environments", IEEE Transactions on Knowledge and Data Engineering (TKDE), p. 1112, vol. 16, (2004). Published
- D. V. Kalashnikov, S. Prabhakar, S. E. Hambrusch, "Main Memory Evaluation of Monitoring Queries Over Moving Objects", Distributed and Parallel Databases: An International Journal, p. 117, vol. 15, (2004). Published
- W. Aref, A. Catlin, A. Elmagarmid, J. Fan, M. Hammad, I. Ilyas, M. Marzouk, S. Prabhakar, Y. Tu, X. Zhu, "VDBMS: A Testbed Facility for Research in Video Database Benchmarking", ACM/Springer Multimedia Systems Journal, p. 575, vol. 9, (2004). Published
- R. Cheng and S. Prabhakar, "Managing Uncertainty in Sensor Databases", SIGMOD Record, p. 41, vol. 32, (2003). Published
- M. J. Atallah, S. Prabhakar, "(Almost) Optimal Parallel Block Access for Range Queries", Information Sciences: An International Journal, p. 21, vol. 157, (2003). Published
- S. Prabhakar, D. Agrawal, A. El Abbadi, "Disk Allocation for Fast Range and Nearest-Neighbor Queries", Distributed and Parallel Databases: An International Journal, p. 1, vol. 14, (2003). Published
- S. Prabhakar, D. Agrawal, A. El Abbadi, "Optimal Scheduling Algorithms for Tertiary Storage", Distributed and Parallel Databases: An International Journal, p. 1, vol. 14, (2003). Published
- S. Prabhakar, D. Agrawal, A. El Abbadi, A. Singh, T. Smith, "Browsing and Placement of Multi-Resolution Images on Parallel Disks", ACM/Springer Multimedia Systems Journal, p. 459, vol. 8, (2003). Published
- Yufei Tao, Reynold Cheng, Xiaokui Xiao, Wang Kay Ngai, Ben Kao, Sunil Prabhakar, "Indexing Multi-Dimensional Uncertain Data with Arbitrary Probability Density", Proc. of the International Conference on Very Large Databases (VLDB), p. 922, vol. , (2005). Published

- Reynold Cheng Ben Kao, Sunil Prabhakar Alan Kwan, Yicheng Tu, "Adaptive Stream Filters for Entity-based Queries with Non-Value Tolerance", Proc. of the International Conference on Very Large Databases (VLDB), p. 37, vol. , (2005). Published
- Reynold Cheng Sarvjeet Singh Sunil Prabhakar, "U-DBMS: A Database System for Managing Constantly-Evolving Data (Demonstration)", Proc. of the International Conference on Very Large Databases (VLDB), p. 1271, vol. , (2005). Published
- Yicheng Tu Mohammed Hefeeda Yuni Xia Sunil Prabhakar, "Control-based Quality Adaptation in Data Stream Management Systems", Proc. of the International Conference and Workshop on Database and Expert Systems Applications (DEXA), p. 746, vol. , (2005). Published
- Yicheng Tu Jingfeng Yan Sunil Prabhakar, "Quality-aware Replication of Multimedia Data", Proc. of the International Conference and Workshop on Database and Expert Systems Applications (DEXA), p. 240, vol. , (2005). Published
- Yuni Xia, Sunil Prabhakar Jianzhong Sun, Shan Lei, "Indexing and Querying Constantly Evolving Data Using Time Series Analysis", Proc. of the 10th International Conference on Database Systems for Advanced Applications (DASFAA 2005),, p. 637, vol. , (2005). Published
- Yuni Xia, Sunil Prabhakar Shan Lei, "Indexing Constantly Changing Data with mean-variance tree", Proc. of the ACM Symposium on Applied Computing (ACM SAC), Mobile Computing and Applications Track, p. 1125, vol. , (2005). Published
- Reynold Cheng Yuni Xia, Sunil Prabhakar Rahul Shah, "Change-Tolerant Indexing for Constantly Evolving Data", Proc. of the International Conference on Data Engineering (ICDE), p. 391, vol. , (2005). Published
- Reynold Cheng Yuni Xia, Sunil Prabhakar Rahul Shah Jeffrey S. Vitter, "Efficient Indexing Methods for Probabilistic Threshold Queries over Uncertain Data", Proc. of the International Conference on Very Large Databases (VLDB), p. 876, vol. , (2004). Published
- Radu Sion M. J. Atallah Sunil Prabhakar, "Resilient Rights Protection for Sensor Streams", Proc. of the International Conference on Very Large Databases (VLDB),, p. 732, vol. , (2004). Published
- Y. Tu, S. Prabhakar, A. K. Elmagarmid, and R. Sion, "QuaSAQ: An Approach to Enabling End-to-End QoS for Multimedia Databases", Proc. of the International Conference on Extending Database Technology (EDBT), p. 694, vol. , (2004). Published
- K. Lou, S. Prabhakar, and K. Ramani, "Content-Based Three-Dimensional Engineering Shape Search", Proc. of the International Conference on Data Engineering (ICDE),, p. 754, vol. , (2004). Published
- Radu Sion M. J. Atallah Sunil Prabhakar, "wmdb.: Rights Protection for Numeric Relational Data (Demonstration)", Proc. of the International Conference on Data Engineering (ICDE), p. 863, vol. , (2004). Published
- Y. Tu, J. Sun, and S. Prabhakar, "Performance Analysis of a Hybrid Media Streaming System", Proc. of the ACM/SPIE Conference on Multimedia Computing and Networking (MMCN), p. 69, vol. , (2004). Published
- K. Lou, S. Jayanti, N. Iyer, Y. Kalyanaraman, Sunil Prabhakar, K. Ramani, "A Reconfigurable 3D Engineering Shape Search System Part II: Database Indexing, Retrieval, and Clustering", Proc. of the ASME DETC 03, 23rd Computers and Information in Engineering (CIE) Conference, p. ??, vol. , (2003). Published
- Yuni Xia, Sunil Prabhakar, "Efficient Moving Object Indexing in Location-aware Services", Proceedings of the 2003 International Workshop on Mobile Distributed Computing (MDC'03), p. 414, vol. , (2003). Published
- Reynold Cheng D. Kalashnikov Sunil Prabhakar, "Evaluating Probabilistic Queries over Imprecise Data", proceedings of the ACM International Conference on Management of Data SIGMOD 2003, p. 551, vol. , (2003). Published
- Radu Sion M. J. Atallah Sunil Prabhakar, "Rights Protection for Relational Data", In proceedings of the ACM International Conference on Management of Data SIGMOD 2003, p. 98, vol. , (2003). Published



- Yuni Xia, Sunil Prabhakar, "Efficient Indexing for Moving Object Databases", Proceedings of the 8th International Conference on Database Systems for Advanced Applications (DASFAA), p. 175, vol. , (2003). Published
- D. V. Kalashnikov, Sunil Prabhakar, "Similarity Join for Low- and High-Dimensional Data", Proceedings of the 8th International Conference on Database Systems for Advanced Applications (DASFAA) 2003, p. 7, vol. , (2003). Published
- Reynold Cheng Sunil Prabhakar D. Kalashnikov, "Querying Imprecise Data in Moving Object Environments", Proceedings of the IEEE International Conference on Data Engineering, p. 723, vol. , (2003). Published
- K. Frikken , M. J. Atallah , Sunil Prabhakar , R. Safavi-Naini, "Optimal Parallel I/O for Range Queries through Replication", Proceedings of the 13th International Conference on Database and Expert Systems Applications - DEXA 2002, p. 669, vol. , (2002). Published
- Elisa Bertino, Moustafa A. Hammad, Walid G. Aref, Ahmed K. Elmagarmid, "An Access Control Model for Video Database Systems", CIKM 2000, p. 336, vol. , (2000). Published
- Walid G. Aref, Ihab F. Ilyas, "SP-GiST: An Extensible Database Index for Supporting Space Partitioning Trees", J. Intell. Inf. Syst., p. 215, vol. 17, (2001). Published
- X. Zhu, X. Wu, A. Elmagarmid, A. Feng & L. Wu, "Video Data Mining: Semantic Indexing and Event Detection from the Association Perspective", IEEE Trans. on Knowledge and Data Engineering, p. 665, vol. 17, (2005). Published
- E. Bertino, A. K. Elmagarmid, and M-S. Hacid, "A Logical Approach to Quality of Service Specification in Video Databases.", Multimedia Tools and Applications, p. 75, vol. 23, (2004). Published
- J. Fan, H. Luo, A. K. Elmagarmid, "Concept-Oriented Indexing of Video Databases: Towards Semantic Sensitive Retrieval and Browsing", IEEE Trans. on Image Processing, p. 974, vol. 13, (2004). Published
- J. Fan , X. Zhu, A.K. Elmagarmid, W.G. Aref, L. Wu, "ClassView: Hierarchical Video Shot Classification, Indexing, and Accessing", IEEE Trans. on Multimedia, p. 70, vol. 6, (2004). Published
- E. Bertino, T. Catarci, A. K. Elmagarmid, and M-S. Hacid, "Quality of Service Specification in Video Databases", IEEE Multimedia, p. 71, vol. 10, (2003). Published
- E. Bertino, J. Fan, E. Ferrari, M-S. Hacid, A. K. Elmagarmid, and X. Zhu, "A Hierarchical Access Control Model for Video Database Systems", ACM Transactions on Information Systems (TOIS), p. 155, vol. 21, (2003). Published
- X. Zhu, J. Fan, A. K. Elmagarmid, X. Wu, "Hierarchical video summarization and content description joint semantic and visual similarity", ACM/Springer Multimedia Systems Journal, p. 31, vol. 9, (2003). Published
- E. Bertino, A. K. Elmagarmid, and M-S. Hacid, "A Database Approach to Quality of Service Specification in Video Databases", Sigmod Record, p. 35, vol. 32, (2003). Published
- Athena Vakali, Mohand-Said Hacid, Ahmed K. Elmagarmid, "MPEG-7 based description schemes for multi-level video content classification", Image Vision Comput., p. 367, vol. 22, (2004). Published
- B. Shafiq, H. Fahmi, S. Baqai, A. Khokhar and A. Ghafoor, "Wireless Network Resource Management for Web-Based Multimedia Document Services", IEEE Communication Magazine, p. 138, vol. 41, (2003). Published
- A. Bashandy, E. K. Chong, and A. Ghafoor, "Generalized Quality of Service Routing Protocol Using Intelligent Resource Allocation Scheme ", IEEE Journal on Selected Areas in Communications, Special Issue on Intelligent Services and Applications in Next Generation Networks, p. 450, vol. 23, (2005). Published
- A. Ghafoor, "Distributed Multimedia Information Systems: An End-to-End Perspective

", Jour. of Multimedia Tools and Applications  
 , p. , vol. , ( ). Accepted

A. Ghafoor, "Distributed Multimedia Information Systems: An End-to-End Perspective  
 ", Jour. of Multimedia Tools and Applications  
 , p. , vol. , (2005). Accepted

R. Bhatti, B. Shafiq, M. Shehab, and A. Ghafoor, "Distributed Access Management in Multimedia IDCs  
 ", IEEE Computer, p. 60, vol. 38, (2005). Published

Sheikh, W.; Shafiq, B.; Paul, R.A.; Ghafoor A., "Provision of multimedia services in a mobile ad hoc network  
 ", Proceedings of the 9th IEEE International Workshop on Object-Oriented Real-Time Dependable Systems, p. 87, vol. , (2003). Published

### 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 Non cooperative 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, Published  
 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, Published  
 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, Published  
Collection: International Conference on Visual Communication and Image Processing  
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.", (2005). Book, Published  
Editor(s): K. Park and W. Willinger  
Collection: The Internet as a Large-Scale Complex Systems.  
Bibliography: Oxford Press
- K. Park and H. Lee, "On the effectiveness of probabilistic packet marking for IP traceback under denial of service attack.", (2001). Conference Publication, Published  
Collection: Proc. IEEE INFOCOM '01, 2001  
Bibliography: Submitted
- H. Ren and K. Park, "Optimal aggregate-flow scheduling.", (2001). Conference Publication, Published  
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, Published  
Collection: Proc. IEEE International Conference on Distributed Computing  
Systems  
Bibliography: 2001

A. Vakali and C. Stupa, "A QoS based Disk Subsystem", (2001). Conference Paper, Published  
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, Published  
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

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, Published  
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, Published  
Bibliography: Proceedings of SPIE ITCOM

Minseok Kwon and Sonia Fahmy, "Topology-Aware Overlay Networks for Group Communication", (2002). Proceedings, Published  
Bibliography: Proceedings of ACM NOSSDAV

Minseok Kwon and Sonia Fahmy, "TCP Increase/Decrease Behavior for Explicit Congestion Notification (ECN)", (2002). Proceedings, Published  
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, Published

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

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

Collection: Encyclopedia of Telecommunications

Bibliography: John Wiley & Sons

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

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, Published

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, Published

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, Published  
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, Published  
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, Published  
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, Published  
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, Published  
Collection: 13th International Conference on Database and Expert Systems Applications  
Bibliography: DEXA
- Jiangtao Li and Sunil Prabhakar, "Data Placement for Tertiary Storage", (2002). Proceedings, Published  
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, Published  
Collection: 7th International Symposium on Spatial and Temporal Databases  
Bibliography: SSTD
- Christopher Clifton, "The Handbook of Data Mining: Security and Privacy.", (2003). Book, Published  
Editor(s): Nong Ye (ed.)  
Collection: The Handbook of Data Mining  
Bibliography: Lawrence Erlbaum Associates, Inc. Mahwah, NJ.
- Christopher Clifton and Jaideep Vaidya, "Privacy Preserving Association Rule Mining in Vertically Partitioned Data", (2002). Conference Proceedings, Published  
Collection: The Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining  
Bibliography: July 23 - 26, 2002, Edmonton, Alberta, Canada
- Christopher Clifton and Murat Kantarcioglu, "Privacy-preserving Distributed Mining of Association Rules on Horizontally Partitioned Data", (2002). Conference Proceedings, Published  
Collection: The ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery (DMKD'2002)  
Bibliography: June 2, 2002, Madison, Wisconsin
- D. V. Kalashnikov, Sunil Prabhakar, W. G. Aref, S. E. Hambrusch, "Efficient Evaluation of Continuous Range Queries on Moving Objects", (2002). Conference Proceedings, Published  
Collection: Proceedings of the 13th International Conference on Database and Expert Systems Applications -  
Bibliography: DEXA 2002, Aix-en-Provence, France
- Walid G. Aref, Ann Christine Catlin, Jianping Fan, Ahmed K. Elmagarmid, Moustafa A. Hammad, Ihab F. Ilyas, Mirette Marzouk, Xingquan Zhu, "A Video Database Management System for Advancing Video Database Research", (2002). Conference Proceedings, Published  
Collection: Proceedings of the Multimedia Information Systems Workshop  
Bibliography: MIS 2002, International Workshop on Management Information Systems. November, 2002. Tempe, Arizona
- Ihab F. Ilyas, Walid G. Aref, Ahmed K. Elmagarmid, "Joining Ranked Inputs in Practice", (2002). Conference Proceedings, Published  
Collection: In Proceedings of the 28th VLDB Conference.  
Bibliography: VLDB Conference. Hong Kong, China. 2002
- Aref, W., El-Bassyouni, K., Kammel, I. and Mokbel, M., "Scalable QoS-aware disk scheduling", (2002). Conference Proceedings, Published



Collection: In IDEAS-02 Proceedings of the International Database Engineering and Applications Symposium  
Bibliography: IDEAS-02 Alberta, Canada. July, 2002

Mohamed F. Mokbel, Walid G. Aref, and Ananth Grama, "Spectral LPM: An Optimal Locality-Preserving Mapping using the Spectral (not Fractal) Order", (2002). Conference Proceedings, Published  
Collection: Proceedings of the 10th International Conference on Geographic Information System  
Bibliography: ACM GIS, McLean, VA, Nov. 2002

Moustafa A. Hammad, Walid G. Aref and Ahmed K. Elmagarmid, "Search-based Buffer Management Policies for Streaming in Continuous Media Servers", (2002). Conference Proceedings, Published  
Collection: In Proceeding s of IEEE International Conference on Multimedia and Expo  
Bibliography: Proceeding s of IEEE Conference on Multimedia and Expo, Aug., 2002, Lausanne, Switzerland

M. Elfeky, V. Verykios, and A. Elmagarmid, "TAILOR: A Record Linkage Toolbox", (2002). Conference Proceedings, Published  
Collection: In Proc. of ICDE 2002, the 18th Int. Conf. on Data Engineering  
Bibliography: ICDE 2002, San Jose, CA, February 2002

M. Elfeky, A. Elmagarmid, and T. Ghanem, "Quality Assurance of Government Databases", (2002). Conference Proceedings, Published  
Collection: 3rd National Conf. on Digital Government Research  
Bibliography: Conf. on Digital Government Research, Los Angeles, CA, May 2002

K. Park, "The Internet as a complex system", (2005). Book, Published  
Editor(s): The Internet as a Large-Scale Complex System  
Collection: K. Park and W. Willinger  
Bibliography: SFI Studies in the Sciences of Complexity, Oxford University Press

### 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:**

The EduMed web-based multimedia database application, a medical education query and retrieval system, 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:**

Our VDBMS 1.0 video database management system is a complete, installable system. The VDBMS 1.0 Windows-based client and Unix-based database server will shortly be made available for download from the web. This will include the code for both client and server, installation instructions, and video data sets.

Our new SP-GIST extensible index structure supports the class of space-partitioning unbalanced trees, and is targeted for use in data mining, GIS, CAD/CAM and other emerging database applications.

We have developed our SP-GIST index code over the Web for use and testing by other colleagues. The code and user manual can be downloaded at [http://www.cs.purdue.edu/homes/aref/dbsystems\\_files/SP-GIST](http://www.cs.purdue.edu/homes/aref/dbsystems_files/SP-GIST).

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 NLNAR 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 are 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.

**Categories for which nothing is reported:**

## **Purdue Multimedia Support Infrastructure - Activities**

The MSI research project addresses the integration of multimedia databases, networking and security – three pillars of today’s digital society – through a computing environment capable of meeting user and application requirements on the Internet. The three subsystems, (1) context-sensitive information retrieval, storage and dissemination, (2) QOS-sensitive information transport, and (3) security-sensitive information processing, are coordinated through the MSI framework to achieve a seamless and efficient single-system image. MSI presents a user-friendly abstraction that facilitates single-stop computing in the digital space.

We describe in the following sections the major MSI project activities for the entire award period.

### **1 Multimedia Databases: Context-sensitive Information Retrieval, Storage and Dissemination**

Real world video-based applications require database technology that is capable of storing digital audio-visual information in the form of video databases and providing content-based video search and retrieval. The development of Purdue University’s VDBMS video database management system within the MSI project was motivated by the requirements of video-based applications to retrieve portions of video data based on content and, more particularly, by the need for testbed facilities to facilitate research in the area of video database management. VDBMS provides a full range of functionality for video as a well-defined data type, with its own description, parameters and applicable methods. The development and integration of a video data type into the database management system achieves a clear separation between the video processing and database components. This allows video-based application design to focus on details of the application itself, while relying on the underlying video framework components for storage, search, retrieval, analysis and presentation of the video data. During the past year, the research problems that were addressed by VDBMS to support the handling of video data included MPEG7 standard multimedia content representation, new algorithms for image-based shot detection, new image processing techniques, investigation of new high-dimensional indexing techniques for accessing the high-dimensional feature vectors extracted by image pre-processing, multimedia query processing and optimization, new query operators, real-time stream management, a search-based buffer management policy, and an access control model for selective, content-based access to streaming video.

Content-based Video Preprocessing. The VDBMS image processing toolkit extracts all low-level features defined by MPEG7 as standard, including color histogram in both HSV and YUV formats, texture tamura, texture edges, color moment and layout, motion and edge histograms, dominant and scalable color, and homogeneous texture. We are currently developing a VDBMS wrapper that

abstracts the extraction, representation, and query of features. This plug-in component allows users to define a new feature, supply its extraction (image processing) algorithm, and query against the feature for image similarity matching. Our wrapper and integration mechanisms incorporate the feature into the query interface, create the schema for database representation and apply the user-provided algorithm during video pre-processing and image-based queries. We provide users with a graphical interface for defining and integrating video segmentation algorithms, feature extraction algorithms, camera motion classification techniques, and other video processing techniques that can be used for content representation and content-based retrieval. This will allow researchers to compare and evaluate alternate methods, improve existing algorithms or develop new ones. We are currently building an interface to support a plug-in component for indexing techniques, so that alternative indexing mechanisms can be implemented, tested, and compared within the VDBMS system.

New VDBMS Query Operators. We have developed a practical, binary, pipelined rank-join query operator NRA-RJ. The operator determines an output global ranking from the input ranked video streams based on a score function. The rank-join algorithm assumes no random access is available on the input streams, and the new VDBMS query operator encapsulates the rank-join algorithm in its *GetNext* operation. The output of NRA-RJ serves as valid input to other operators in the query pipeline, supporting a hierarchy of join operations and integrating easily into the query processing engine of any database system. A modular interface for the integration of new query operators into the VDBMS query processor is currently underway. The interface will support the integration of user-developed operators into the query execution plan, and will also support the performance evaluation and comparison of alternative algorithms for implementing query operators by allowing developers to identify performance metrics and test point locations for collecting measurements and statistics.

Search-based Buffer Management. We have developed an efficient buffer management policy that uses feedback from the search engine to make more accurate replacement and pre-fetching decisions. Top-ranked query results from the query processor are used to predict future video streaming requests, and a weight function determines candidates for caching. By integrating knowledge from the query and streaming components, VDBMS can achieve better caching of media streams, thus minimizing initial latency and reducing disk I/O. In our search-based replacement policy, pages in the buffer pool that are referenced by either current or expected streams are considered for caching. Lookup tables contain pointers to expected streams, which are collected from the search results and checked by the stream manager for matches when determining pages to replace. The stream manager tracks the utilization of the streaming period, and utilizes any fraction of the streaming period unused by current streams to pre-fetch the first segment of the top ranked expected streams into the memory buffer. We evaluated the performance of the search-based by investigating the effects of buffer management on the number of I/Os when referencing the first segment of a requested stream. Experimental results show that initial latency of

the search-based policy is reduced on the average by 20% when compared with traditional policies.

Tertiary storage in multimedia database systems. Tertiary storage in computer systems is supported by a group of low-cost storage technologies. However, the performance of tertiary storage devices is inferior to that of magnetic disks, especially with respect to the response time. For this reason, tertiary storage (TS) has been used only for data backup and archiving in database systems. In our research, we have investigated the issues of storing active data in tertiary storage devices in multimedia database systems. The main idea was to augment the current storage manager in the DBMS to handle data stored in tertiary storage. A magnetic disk partition was dedicated to caching frequently-accessed items in tertiary storage. An add-on software module called the Tertiary Storage Manager (TSM) is responsible for disk cache management and data retrieval from TS devices. The TSM module was implemented on the Video Database Management System (VDBMS) platform. The TS device we used was a 200-disc DVD/CD changer (Model 0200, details can be found in <http://www.dvdchanger.com>) from Powerfile, Inc. The data stored in TS are medical videos in MPEG-i format. Our experiments showed that the natural latency of streaming video from TS was on the order of tens of seconds (10—30 seconds) and the transfer rate can be as high as 16Mbps. The relationship between transfer rate and transfer block size was also investigated. With the introduction of smart pre-fetching and caching algorithms, it is feasible to use TS to support data-intensive applications such as media streaming.

Data placement, migration, and indexing techniques for video and multi-dimensional spatio-temporal data. Several activities have been conducted in this area:

1) Data management schemes for very large amounts of video on hierarchical storage. We have developed a novel caching scheme for secondary storage that when coupled with replication on tertiary storage yields significant reductions in start-up latency for continuous multimedia objects such as video. We have also developed a new placement scheme for tertiary storage that takes into account relationships between objects to reduce expensive swapping of media. Current evaluation of above schemes has been based upon simulation. We are currently experimenting with tertiary devices (Powerfile CD changers) on which to test these schemes. The devices are being incorporated into a prototype multimedia database system being developed. A set of novel placement schemes for tertiary storage that take relationships between objects into account were developed.

2) We have investigated data placement schemes for the efficient retrieval of multi-resolution video from disks. Alternative schemes have been developed and tested using a simulation setup that has been developed on top of available disk simulators.

3) We have developed two indexing techniques for spatio-temporal data to efficiently process large numbers of concurrent, ongoing queries over moving objects: Query-Indexing and Velocity-Constrained Indexing. Query indexing

reverses the role of queries and data for moving objects thereby eliminating the need for frequent updates to an index over data points. Velocity-Constrained indexing exploits a pessimistically high limit on the speed of an object in order to avoid updates to the index. Together the two index structures provide efficient and scalable support for continuous queries over constantly updated data. Main memory evaluation has also been investigated resulting in the development of a novel indexing method (Memory Query Indexing). Using a simple grid based index we are able to achieve an order of magnitude better performance for continuous queries over moving objects than traditional indexes in main memory.

4) Uncertainty Management for Spatio-temporal and sensor data. We have investigated the problem of inherent uncertainty of current data for moving objects and sensor applications. We have studied the management of uncertain data in relational database and the notion of probabilistic queries for constantly evolving data. We have developed a model for data uncertainty and associated probabilistic query evaluation algorithms. We have also introduced the notion of query quality for these probabilistic queries and proposed a number of metrics for quantifying the quality. For efficient and scalable execution, we have investigated the problem of indexing uncertain data in order to improve query evaluation performance. Based upon these developments, a prototype database management system with native support for uncertain data and probabilistic queries has been developed. It is implemented as an extension to PostgreSQL and was demonstrated at VLDB 2005. This prototype system will be made available to the public.

5) A number of new indexing methods have been developed for high update environments, including moving objects. These solutions include Change-Tolerant R-trees, Q+R-trees, the CNG tree, and the Mean-Variance tree. These methods exploit various properties of constantly-changing data in order to provide improved update performance without sacrificing query performance. These structures are based upon topology (e.g. buildings, highways, homes, etc.) and human behavior. Change-Tolerant indexes explicitly optimize for both update and query (as opposed to only queries) and can be tailored to the relative rate of updates to queries.

## **2 Networking: QoS-sensitive Information Transport**

QoS Support in Multimedia Databases. In spite of the fact that research in multimedia databases has covered many key issues, little attention has been devoted to the integration of QoS support at the user level. Due to the heterogeneity in users' preference and the computing power of the client machines, we believe the QoS requirements of each media object access to the database are not identical. Current multimedia database system performs QoS-provisioning, which leads to waste of resources and degradation of QoS for all users under high system contention. We are investigating a QoS-aware multimedia DBMS that dynamically determines media delivery plans according to the system contention level and user requirements. The central part of the DBMS is a Quality-of-Service-aware Query Processor (QuaSAQ) that can be regarded



as an extension to the current query processing and optimization module in DBMSs. In our system, user-level QoS parameters are translated into system QoS (by QoS profiling strategies) and become an augmented component of the query. QuaSAQ then generates alternative plans for media delivery. These plans are evaluated using a cost model and the best one is chosen to be the execution plan. Before execution, QuaSAQ also takes care of resource reservation to satisfy the needs of the chosen plan. We also proposed a novel cost model that quantifies the resource utilization of each plan. In this model, each resource is priced according to its demand and supply. The cost of a plan is given by the total price of the resources involved. The implementation of QuaSAQ is underway.

Purdue Infobahn QoS Testbed. Purdue Infobahn is a QoS testbed for implementing and evaluating the SBS wide area network QoS architecture developed at the Network Systems Lab. It is a private IP-over-SONET backbone network consisting of a number of Cisco 7206 VXR routers whose router software is being customized to implement the SBS IP switching algorithm. The aim of Purdue Infobahn is to benchmark the performance of SBS at delivering user-specified QoS without resorting to resource reservation and admission control, two central mechanisms that impede scalability. Guaranteed services are provided as part of the Purdue Infobahn service suite for applications requiring strict per-flow QoS protection. SBS will be compared with other IETF diff-serv related architectures and benchmarked with multimedia applications in the context of MSI. The Scalable Internet QoS Architecture Project is investigating new solutions to effective QoS provisions on the internet. Issues include the non-cooperative network environment with selfish applications or users, scalable QoS-sensitive services, stratified best-effort QoS provision without resource reservation and provision of stable, adaptable services. Current activities involve the analysis of non-cooperative QoS provision games, the design/analysis/implementation of LAN-scale systems, the design/analysis/implementation of WAN-scale (many-switch) systems and game-theoretic analysis of congestion control protocols. The Self-similar and Multimedia Traffic Control Project is investigating self-similar burstiness, transport of multimedia data with real-time QoS constraints using adaptive redundancy control, game theoretic considerations and fairness, optimality and efficiency.

Minimizing Power Use while Maintaining QoS Guarantees .Predicting Movement in Wireless Communication. Node movement can be exploited to reduce the energy consumption of wireless network communication. Our strategy consists in delaying communication until a mobile node moves close to its *target* peer node, within an application-imposed deadline. We have evaluated the performance of various heuristics that, based on the movement history of the mobile node, estimate an optimal time (in the sense of least energy use) of communication subject to the delay constraint. We also evaluated the impact of the node movement model, length of movement history maintained, allowable delay, single hop versus multiple hop communication, and size of data transfer on the energy consumption. We then analyzed measurement results on an iPAQ pocket PC to quantify energy consumption in executing the prediction algorithms. Our results

showed that, with relatively simple and hence efficient prediction heuristics, energy savings in communication can significantly outweigh the energy expenses in executing the prediction algorithms. Moreover, we showed that it is possible to achieve robust system performance across diverse node movement models.

Querying Imprecise Data in Moving Object Environments. We are investigating the imprecision of the positions of moving objects. Given limited network bandwidth, it is impossible for a database server to record the exact location of moving objects. Querying on the database values can result in error, but if the uncertainty is limited, we can bound the uncertainty region of each moving object and query the uncertainty region instead. We have studied how to augment probability values to queries in moving objects. We are currently investigating algorithms for evaluating range queries and nearest-neighbor queries, and further derived algorithms for objects moving in two patterns: line-moving and circle uncertainty. We have conducted experiments to study the trade-off between efficiency and precision of the algorithms.

This work has been extended to encompass the study of probabilistic querying of general sensor values which are constantly changing. Based on the type of the probabilistic results, we classified queries according to whether they are aggregate queries and whether the results are entity or value-based. We then studied algorithms for different kinds of queries, and further defined the quality of a query result. We studied and evaluated several update policies that tried to improve the quality of result for a query. We continue to experimentally evaluate more probabilistic queries and study the practical issues of probabilistic queries. Our most recent variant of probabilistic queries allows the user to specify a probability threshold in the query, and only the objects with probability values higher than the threshold are reported. Several methods can reduce the computation effort required by the original algorithm in this scenario. We plan to evaluate the effectiveness of these methods through experimentation.

### **3 Security: Security-sensitive Information Processing**

ATM Network Security. In this project, we are investigating the problem of providing effective security services in wide area networks which incorporates the overhead impact on QoS. Our architecture is reactive in nature, minimizing the services provided proactively due to their high operational cost. We have implemented a prototype architecture based on SNMP which adaptively manages security services by detecting and responding to anomalies using a form of monitoring-processing-control. Our activities include the design, implementation and analysis of ATM network security architectures. We are addressing the following issues: (1) an ATM network security model, (2) native ATM solutions (telephony), (3) IP-over-ATM solutions (data communication), and (4) integration of ATM QoS/efficiency with security.

Watermarking. Much of the current online content can be found under the form of structured aggregates of multiple types of data. These structures can be represented abstractly as graphs, usually in a natural way, otherwise via a suitable mapping to graphs, and are characterized by value lying both in the

structure itself (the aggregation graph) and in the (possibly non-structured) content (the nodes). Examples include XML documents, complex Web Content, relational DBMS data, VRML and similar environmental representations, structured financial and B2B interaction data, workflow and planning descriptions. My dissertation explores the issue of securing such content, more specifically protecting rights through Information Hiding and detecting intrusions in systems with underlying structured data sources.

Rights protection in the digital framework has become particularly relevant with the advent of the Internet, the Web and the associated ability to distribute online valuable content. Different avenues for rights protection are available, each with its own advantages and drawbacks. Enforcement by legal means is usually ineffective in preventing theft of copyrighted works, unless augmented by a digital counter-part, for example watermarking. Legal and technological approaches need to work hand in hand to achieve such protection. The watermarking framework relies on the idea of hiding rights association identifiers within the digital valuable, enabling ulterior court- time proofs of rights ownership by recovery of the hidden information. Most previous work on rights protection through watermarking has consisted of information hiding techniques for specific types of media: images, audio and video. Given the above, we recognized a genuine need for a complete formal understanding of rights protection in the generic framework of structured content, as opposed to attempts within a given data-specific sub-domain. This research has led to the development of a foundational framework for generic watermarking for rights protection and the analysis of important theoretical bounds. We developed and implemented watermarking algorithms for rights-protection of semi-structured data. One particular type of structure is exhibited by relational data through associated database and semantic constraints. Protecting rights over relational data is of ever increasing interest, especially considering areas where sensitive, valuable content is to be outsourced. A good example is a data mining application (e.g. supermarket sales database, oil drilling data, financial data), where a data set is sold in pieces to parties specialized in mining it. Thus, our research focused also on this closely related area in rights protection and resulted in the first industry-level database watermarking application (<http://www.cs.purdue.edu/homes/sion/wm/wmdb>), which deploys information hiding techniques to solve the above issues.

In the broader framework of access control security, we also designed algorithms to detect network intrusions by training for normality in data access patterns in a system with an underlying hyper-linked structure. An instance of this is the detection of intrusions in password protected web portals.

## **4 Additional External Funding**

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 (SF1) and National Science Foundation to hold a joint SF1/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.

## 5 Specific Activities

In this section, we highlight specific activities in each of the three major areas of the project: multimedia databases and storage management systems, networking, and security.

### 5.1 *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 nearoptimal 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 AI 000 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, pre-fetching, 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.

## **5.2 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 nsbased 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-overSONET 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 'A/LAN 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.

### **5.3 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 bench marked 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.

## 6 Additional Research Funds and Grants

We have been successful in obtaining research funds and grants for related research from several sources including:

1. **Kihong Park**. The Internet as a Complex System, Santa Fe Institute, Fellow-at-Large, PI, 08/00-07/01, \$5,000.
2. **Kihong Park**. Toward Scalable Solutions for Distributed Denial of Service Attack Prevention, DARPA, PI, 7/01-6/03, \$442,284.
3. **Kihong Park**. QoS-Aware, Secure and Reliable Distributed Scheduling, Xerox, PI, 11/00-10/03, \$45,000.
4. **Kihong Park**. Network Security Management Framework Research, ETRI, PI, 11/02-12/05, \$207,265.
5. **Kihong Park**. Predictable, Scalable QoS Routing for Ad Hoc Wireless Networks Based on Heavy-tailed Statistics, IAI/Army RDECOM, PI, 08/05-02/06, \$35,000.
6. **Kihong Park** and Jennifer Hou. ITR: Multiple Time Scale Traffic Control for Next Generation Internets, National Science Foundation, PI, 09/00-09/03, \$460,370.
7. **Kihong Park** and Walter Willinger. The Internet as a Large-Scale Complex System, National Science Foundation, PI, 03/01-02/02, \$15,000.
8. Ann Hendrich (Ascension Health), Marilyn Chow (Kaiser Permanente), Nelson Lee (Rapid Modeling), William Cleveland, **Chris Clifton**, Jason Abrevaya, A Multi-Site Study of How Medical Surgical Nurses Spend Their Time: A Baseline Study in Preparation for an Electronic Health Record and an Evidenced-based Nursing Unit Design, Robert Wood Johnson Foundation, April 2005-March 2006. (Purdue CS portion \$59,481).
9. Clement J. McDonald (PI) et al., **Chris Clifton** (Purdue PI) A Center of Excellence in Medical Informatics to Provide an Advanced Infrastructure for Human Research: A Catalyst for Indiana Research, Indiana 21st Century Fund, August 2004-August 2006, \$3,832,196 (Lead: Regenstrief Institute for Healthcare; Purdue portion \$50,000.)

10. **Chris Clifton**, AnHai Doan, **Ahmed Elmagarmid**, Gunther Schadow, Dan Suci, Privacy-Preserving Data Integration and Sharing NSF ITR, September 2004-August 2007, \$1,000,000.
11. **Chris Clifton** and Mikhail Atallah, Collaborative Research: Distributed Data Mining to Protect Information Privacy, NSF Small ITR, August 2003 - August 2006, \$276,274. (Collaborative with Wenliang (Kevin) Du, Syracuse University - funded under separate award.)
12. **Chris Clifton, Walid Aref, Ahmed Elmagarmid, Arif Ghafoor, and Sunil Prabhakar**, Knowledge Projection for TeleMaintenance, Naval Surface Warfare Center, Crane Division / EG&G October 21, 2002 - September 30, 2005, \$1,746,813.
13. **Christopher Clifton, Ahmed K. Elmagarmid**, Transportation, Distribution and Logistics: a Strategic Opportunity for Indiana and Purdue. Central Indiana Corporate Partnership. 7/7/2003-8/1/2005. \$285,000.
14. Susanne E. Hambruch, **Sunil K. Prabhakar, Walid Aref**, Query Processing in Pervasive Location Aware Computing Environments. National Science Foundation. 9/1/2001-8/31/2005. \$226,000.
15. **Walid Aref**, Research and Development of Database Technologies for Modern Applications (Career Award). National Science Foundation. 9/15/2001-9/30/2006. \$300,000.
16. **Walid Aref, Ahmed K. Elmagarmid, Aref Ghafoor**, A Test-bed Facility for Research in Video Database Benchmarking. National Science Foundation. 8/15/2002-8/31/2005. \$135,856.
17. **Sunil K. Prabhakar**, Efficient I/O for Modern Database Applications National Science Foundation (Career Award). 10/1/2000-9/30/2005. \$240,000.

# **Purdue Multimedia Support Infrastructure - Findings**

The Purdue Multimedia Support Infrastructure project has built the necessary networking and database infrastructure to support experiments on end-to-end delivery of multimedia data. During the five years of this project, we have designed and built several prototype systems to experiment with and demonstrate tools and algorithms for the three subsystems of the MSI infrastructure: (1) context-sensitive information retrieval, storage and dissemination, (2) QOS-sensitive information transport, and (3) security-sensitive information processing.

This document describes our major findings for the entire period of the award.

## **1 Multimedia Databases and Storage Management Systems**

### ***1.1 EduMed: A Web-based Multimedia Databases for e-Learning in Medical Education***

EduMed is a prototype of a medical education system to support rapid classroom access to online, interactive multimedia medical materials for teaching and learning. The system handles ad hoc video queries based on keywords assigned to video clips, returning representative (key) frames to identify selected clips, and providing quality streaming playback for retrieved video clips.

The EduMed project has brought together researchers, medical practitioners, educators, service providers, and product developers to create an infrastructure for the research, development and delivery of multimedia medical data that could support (1) management and use of multimedia data, (2) annotation of medical video data with text and audio commentaries by physicians, (3) content-based search and retrieval of video data using indices built from the annotations, (4) image-based search and retrieval of video data using meta-data generated by image processing techniques, and (5) coordinated storage and retrieval of video data from remote sites.

The development of this prototype was based on four separate activities: (1) the technology transfer of advanced multimedia functionality to the database engine, TITANIUM, marketed Micro Data Base Systems, Inc., (2) joint efforts with Indiana University Centers of Medical Education to produce relevant, high quality medical education videos, compress them into MPEG format and apply semantic- and image-based processing to create appropriate medical meta-data to describe video clip content, (3) collaborations with faculty from the IU School of Medicine to determine user requirements and functional specifications, and (4) multidisciplinary applications-oriented research and development at Purdue University to create a viable prototype. Cooperative efforts with the Technology Transfer Initiative (TTi) of the Krannert Graduate School of Management produced a commercialization plan for EduMed.

## **1.2 OpenSesame – Real Time Content Distribution**

The most recent quality-of-service tool, OpenSesame, is a legacy-compatible scalable real-time content distribution system that allows turnkey, high-performance, user-sensitive multimedia streaming. Presently, OpenSesame is available to selected members of the Purdue Computer Science Department who use it to receive cable TV on their desktop, laptop, or handheld computer.

OpenSesame allows resource contention algorithms to be benchmarked and evaluated with real users, continuing in the broadband wireless arena where Berkeley's influential INDEX project left off. Resources under contention include shared cable channels, bandwidth, and QoS. System-wide stability, efficiency, and user satisfaction are being studied under a range of policies, including market-oriented priority control.

## **1.3 The Nile Stream Database System**

We have developed Nile, a stream database system that allows for the processing of continuous and snap-shot queries over data streams. Nile extends the query processor engine of a prototype object-relational database management system, to process continuous queries over data streams. Nile supports extended SQL operators that handle sliding-window execution as an approach to restrict the size of the stored state in operators such as join. More specifically, Nile supports the following features: (1) Efficient and correct pipelined execution of sliding window queries over multiple data streams. The correct execution is enforced by two novel pipelined scheduling approaches: the Time Probing approach and the Negative Tuple approach. (2) Scalability in terms of the number of queries and the number of data streams. (3) Access control to accept/register new continuous queries and new streams. (4) Providing guarantees for Quality of Service and Quality of Answers. (5) Online stream summary manager. (6) Integrating online data mining tools in query processing over data streams. (7) Approximate window join processing and joining in a network of data streams.

Nile uses several approaches to schedule the pipelined execution of sliding window operators. The first approach (Time Probing) uses window-based operators; the second approach (Negative Tuple) uses a special operator, W-Expire, to emulate tuple addition to and subtraction from the window content. In addition, the Negative Tuple approach alleviates the operators from checking the window constraint. The third approach (Hybrid) approach uses a mix of the first two approaches. Nile introduces several window operators, besides window join, such as sliding window DISTINCT, Aggregate, and Set operations. In addition, Nile implements each operator as a separate preemptive system-scheduled thread. The operators communicate with each other through a network of FIFO queues.

Nile provides flexible interface to define new streaming sources and execute sliding window queries. Based on the Nile system, we have integrated retail applications, spatio-temporal applications, and video security systems. The

applications use real data sets that include retail transactions from Wal-Mart stores, spatio-temporal data, and video data streams.

### ***1.4 Online Incremental periodic patterns discovery in real-time data streams***

We have developed the first online incremental algorithm to discover periodic patterns in real-time data streams. Real-time data streams occur in a variety of situations, including telecommunications and video surveillance networks. Our algorithm allows, for example, a telephone company to analyze telephone calling patterns and predict periods of high and low traffic so that resources can be allocated efficiently. We have completed a number of experiments with both synthetic and “real” data to perform an empirical analysis of our algorithm. These experiments confirm that the algorithm is correct, discovers periodic patterns in a practical manner, and shows that the patterns are highly accurate.

### ***1.5 VDBMS - Video Database Management System***

**VDBMS 1.0 is a complete** video database management system, offering a Windows- **based** query interface client and a **Unix-based DBMS server**. **The system has been** presented and demonstrated at numerous conferences, and we are currently preparing the system to be placed in the public domain for download and installation. The downloaded system will include the complete source code, installation instructions, test video data sets and user support. The public domain VDBMS 1.0 (built on top of the Shore storage manager and Predator object relational database manager) includes the following advanced video database technology developed by the VDBMS research group within the MSI project: (1) video as a native database data type, (2) video- preprocessing and feature extraction for image-based content representation and context-sensitive video query and retrieval, (3) GiST indexing for high-dimensional feature-based search, (4) a real-time video stream manager to provide continuous media streaming, (5) a search-based buffer management policy (pre-fetching and replacement) for improved streaming performance, (6) a rank-join query operator for multi-feature image queries, and (7) video processing during streaming which provides image-area blurring for content-based secure access control.

While investigating, developing, and testing the fundamental components required to support full video database functionality in VDBMS 1.0, we also utilized our system as a testbed for integrating and evaluating video processing technologies from other sources. As such, the system has provided an environment for testing the correctness and scope of algorithms, measuring the performance of algorithms in a standardized way, and comparing the performance of different implementations of a component. The VDBMS project extended this concept by constructing video component wrappers with well-defined interfaces that allow components to be easily modified or replaced, and then developed corresponding semi-automatic mechanisms for integrating these components into VDBMS. The resulting VDBMS system is a flexible, extensible

framework that can be used by the research community for developing, testing and benchmarking video database technologies.

We have also developed VDMBS 1.1, a Linux-based DBMS server with a Windows-based query interface client. VDBMS 1.1 is a generalization of the VDBMS concept for the *video* data type; our new system has been designed to handle the *stream* data type (e.g., sensor readings, online transactions and network traffic.) The capabilities of VDBMS 1.1 for stream support developed within the MSI project are: (1) a new stream manager which operates as an interface between outside stream devices and internal processing, (2) a stream scan, (3) operators for continuous queries, such as the VDBMS window-join, and (4) support for multiple continuous query optimization and execution.

## **1.6 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 highdimensional 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.

### **1.7 Semantic Video Databases**

Several content-based video retrieval systems have been proposed in the past, but they still suffer from the following challenging problems: semantic gap, semantic video concept modeling, semantic video classification, and concept-oriented video database indexing and access. We propose a framework that includes (1) semantic-sensitive video content representation, by using principal video shots to enhance the quality of features, (2) semantic video concept interpretation, by using flexible mixture model to bridge the semantic gap, (3) a novel semantic video-classifier training framework, by integrating feature selection, parameter estimation, and model selection seamlessly in a single algorithm; and (4) a concept-oriented video database organization technique through a domain-dependent concept hierarchy to enable semantic-sensitive video retrieval and browsing.

### **1.8 Video Mining**

To achieve more efficient video indexing and access, we introduce strategies for video content structure and events mining. The video shot segmentation and representative frame selection strategy are first utilized to parse the continuous video stream into physical units. Video shot grouping, group merging, and scene clustering schemes are then proposed to organize the video shots into a hierarchical structure using clustered scenes, scenes, groups, and shots, in increasing granularity from top to bottom. Then, audio and video processing techniques are integrated to mine event information, such as dialog, presentation and clinical operation, from the detected scenes. Finally, the acquired video content structure and events are integrated to construct a scalable video skimming tool which can be used to visualize the video content hierarchy and event information for efficient access.

### **1.9 Semantic Grouping of Video Shots**

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.

### **1.10 Detection and Tracking**

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



## **1.11 Indexing Video Streams**

We have developed advanced video database indexing technologies for high-dimensional feature-based search and real-time video streaming. A new extensible index structure, termed SP-GiST, is proposed to support a class of data structures, mainly the class of space partitioning un-balanced trees. Simple method implementations demonstrate how SP-GiST can behave as a k-D tree, a trie, a quadtree, or any of their variants. Issues related to clustering tree nodes into pages as well as concurrency control for SP-GiST have been addressed. A dynamic minimum-height clustering technique is applied to minimize disk accesses and to make using such trees in database systems possible and efficient. A prototype of SP-GiST has been implemented.

## **1.12 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.

*MAQP A multimedia authoring, querying, and presentation tool.* 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 web-enabled 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.

### **1.13 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 reevaluate 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.

### **1.14 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.

### ***1.15 Data placement, migration, and indexing techniques***

Major findings in the context of data placement, migration, and indexing for video and multi-dimensional spatio-temporal data are as follows:

- a) Prefix caching and replication. We have developed a novel caching scheme for secondary storage that when coupled with replication on tertiary storage yields significant reductions in start-up latency for continuous multimedia objects such as video.
- b) Edge-merge placement. We have developed a new placement schemes for tertiary storage that takes into account relationships between objects to reduce expensive swapping of media.
- c) Query-Indexing and Velocity-Constrained Indexing. We have developed two new indexing techniques for spatio-temporal data to efficiently process large numbers of concurrent, ongoing queries over moving objects. These are: Query Indexing and Velocity Constrained Indexing.
- d) Broadcast techniques for multi-dimensional data and indexes. Techniques for broadcasting indexed multi-dimensional data to save battery time as well as query time have been developed.

- e) Memory Query Indexing using a simple grid based index we are able to achieve an order of magnitude better performance for continuous queries over moving objects than traditional indexes in main memory.
- f) Watermarking algorithms for numeric sets have been developed. A watermarking scheme for relational databases has been developed. This is the only scheme that satisfies the important requirement of preserving user-specific notions of data quality. A generalized notion of acceptable watermarking has been developed for the case of structured data. Watermarking techniques for XML are being investigated. Extensions of the relational watermarking scheme to non-numeric data are currently under submission.
- g) Algorithms for evaluating range queries with data replicated on disks. An almost optimal placement algorithm with replication was developed -- the first published results on replicated placement for this problem.
- h) Main memory indexing techniques for evaluating spatial joins have been developed. In addition, the importance of choosing the parameters for evaluation has been pointed out.
- i) Novel index structures for moving objects have been developed. These structures are based upon topology (e.g. buildings, highways, homes, etc.) and human behavior. More recently, we have developed Change-Tolerant indexes that are designed for high data update rate environments. These indexes explicitly optimize for both update and query (as opposed to only queries) and can be tailored to the relative rate of updates to queries.
- j) A novel notion of probabilistic queries for moving objects and sensor databases has been developed. Algorithms for computing answers to these probabilistic queries have been developed. The notion of query quality for probabilistic queries has been developed. A classification of probabilistic queries and associated quality metrics has been presented. This is the first work to investigate these issues.
- k) We have developed novel index structures for uncertain data represented as intervals and associated probability distribution functions.

## **2 Networking**

### ***2.1 QoS under Diverse Network Conditions and Resource Constraints to assure Quality of Presentation***

Emerging Internet applications involve accessing large volume of multimedia information from remote servers and databases. Transmission of multimedia data over a broadband network requires special support from the underlying routers and switches. The objective of the network support is to guarantee the quality-of-presentation (QoP) required by the multimedia client(s) at the destination(s). To achieve this objective, the network must provide guaranteed quality-of-service (QoS) under diverse network conditions and resource constraints. As a part of this research, we have proposed a general formulation for the problem of QoS routing with resource allocation for data networks. The formulation incorporates the hardware and software implementation and exhibits three important features.

First, it takes into consideration the relationship between the QoS metric and the allocated resources. Second, it allows independent allocation of different resources coupled with routing. Third it captures the interrelationship among various QoS metrics, including jitter delay, data droppage ratio, average bandwidth, and end-to-end delay. Based on the proposed formulation, we have developed a dynamic programming algorithm to find a route between the source and destination nodes and determine the amount of resources along the intermediate nodes to satisfy the QoS requirements.

## ***2.2 Networking: QOS-sensitive information transport***

The Quality-of-Service-aware Query Processor (QuaSAQ) system extends query processing and optimization in the DBMS by dynamically establishing a media delivery plan based on system contention levels and user requirements. User-level QoS parameters are translated into system QoS requirements which become an augmented component of the query. QuaSAQ generates alternative plans for media delivery, which are evaluated using a cost model. QuaSAQ handles resource reservation to satisfy the needs of the chosen plan.

As part of the Purdue Infobahn Testbed Projects, the Scalable QoS Provision Architecture is a scalable, differentiated services 'A/AN architecture for QoS-sensitive applications with elastic requirements using aggregated flow control and class-based label switching. This work provides a theory of aggregate-flow scheduling which complements the well-known per-flow theories. Akin to real-time scheduling theory showing the existence of optimal real-time schedulers (e.g., EDF, RMS), aggregate-flow theory shows that there are optimal aggregate-flow schedulers which provide the theoretical foundation for differentiated services. The Multimedia and Self-Similar Traffic Control project designs next generation end-to-end Internet protocols for QoS-sensitive transport of real-time - MPEG I & II video/audio - interactive, and bulk data traffic (file transfers). The traffic controls which follow the multiple time scale traffic control framework achieve significant performance improvement by exploiting large time scale predictability structure present in self-similar and heavy-tailed Internet workloads. The multiple time scale traffic control framework is able to mitigate the reactive cost of feedback traffic controls in broadband wide area networks characterized by a high delay- bandwidth product by engaging predictability information exceeding the time scale of the round-trip time. These protocols have been implemented both for TCP and UDP, and run over UNIX and Windows NT.

## ***2.3 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 Test bed 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.

## **2.4 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 queuing framework -m-class  $G/G/i$  queuing 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 queuing -which uses the conservation law (a dot product functional which applies to all work-conserving schedulers and their queuing 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  $M/G/1$ ,  $G/M/i$ , 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 facto 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.lib access points, and a number of Compaq iPAQ Pocket PC H3850 handhelds running our QoS environment over/inside Windows XP and Windows CE.

## **3 Security**

### ***3.1 Security-sensitive information processing***

The Adaptive Network Security and Fault Tolerance Project has integrated network security services with efficiency and QoS such that processing overhead is effectively managed and a user-specified trade-off is achieved. Critical security services (e.g., confidentiality, authentication) - user programmable - are proactively affected, while other user-selected services (e.g., intrusion detection, DoS monitoring) are reactively handled. AdSec is a prototype system built on top of SNMP using distributed agent technology which combines proactive and reactive service protection. More recent work includes scalable solutions to distributed denial of service attack prevention based on a novel approach called route-based distributed packet filtering.

### ***3.2 Access Control for Video Databases***

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 trace back 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.

### ***3.3 Proactive and Scalable Network Architecture for distributed DoS***

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 105 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.