The MSI research project continues to address 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. MSI project activities for the past year are the following:

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 exiting 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 C200, details can be found in http://www.dvdchanger.com) from Powerfile, Inc. The data stored in TS are medical videos in MPEG-1 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.

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.

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.

Equipment and payroll for the project period 2002-2003

Equipment	$54,183
Includes Optiplex GX260 Small Minitower with Pentium REG4 Processor (qty 12), Dell Precision TM Workstation, Dell Optiplex GX260 memory modules (qty 7), 3 Com SuperStack 3 Switch 3300XM 24-Port 10/100 (qty 3), MT4676 Princeton Arcadia AR 3.2 FTX 34'' presentation monitor, 181.6GB 7200RPM Fibre Channel (qty 3), HP Laser Jet 2200DN printer (qty 2), Logitech Quick Cam Express/Pro/Notebook, IBM UltraPort Camera II, miscellaneous software and software agreement/licenses.	

Payroll: Graduate Students, Admin.	$70,000

4/2001-4/2002
Equipment purchases to date, summarized by project area, include:

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

We have been extremely successful in obtaining research funds and grants from numerous sources. Noted among them is an award from the State of Indiana for supporting faculty and graduate students to conduct research and experimental development of a multimedia system for the telemedicine application. These funds are in direct support of this effort. The total funding from the State is over $1.7 million for two years.

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

We have also received internal funding from the Center for Education and Research in Information Assurance and Security (CERIAS) to support one post doctoral fellow, a visiting scholar, and two graduate research assistants. The CERIAS center at Purdue University, which promotes information security research and education, was started by a multi-million dollar grant from the Eli Lilly Foundation. The leadership for this effort is provided by Gene Spafford, its director, who is also a co-PI. Gene Spafford has also received separate NSF funding for his security research.

Our group has received two large grants, one from NCR/Wal-Mart and the other from Hewlett Packard (HP). The NCR grant is for a 1.6 terabyte parallel database engine valued at approximately $7 million. This system not only augments the other storage and compute servers we are purchasing through the NSF grant, but also includes a large sample dataset for warehousing and mining experiments. This parallel server has 80 processors and 400 disk drives. The grant from HP allows us to have 20 workstations and one server to use as desktops for the researchers in this project.

We have received a National Science Foundation ITR grant (joint with Ohio State University and the University of Illinois, who are subcontractors) for studying multiple time scale traffic control aspects of the network research component, a project led by Park (co-PI).

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

We received funding from the Santa Fe Institute (SFI) and National Science Foundation to hold a joint SFI/NSF Workshop "The Internet as a Large-Scale Complex System" (chair: Kihong Park, co-chair: Walter Willinger), March 29-31, 2001. More information is available at http://discuss.santafe.edu/internetcs.

Sunil Prabhakar (senior personnel) has received an NSF Career award to study efficient I/O techniques for a broad range of database applications. The project leverages the equipment that is being acquired through the MSI project. The co-PIs and personnel include several NSF CAREER awardees and a Presidential Young Investigator awardee.

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

Multimedia Databases and Storage Management Systems

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

1. *Enabling content-based access to video data*. We are pursuing activities aimed at developing a comprehensive video database management system that supports content-based retrieval of large-scale video data as a critical component of multimedia document management system. We are developing techniques and tools for content-based retrieval, semantic browsing, and querying of video, including techniques for (1) partitioning video source into meaningful segments to support more effective video representation and indexing, (2) grouping the physical segments into semantically richer units, (3) detecting and recognizing faces, and (4) developing metadata to support fast and effective search and browsing techniques for large-scale video. We have incorporated the MPEG-7 representation of multimedia content into our system, which has significantly affected our schema representation and query management. We are developing a query manager that parses and executes MPEG-7 queries. We are also developing an MPEG-7 wrapper to import pre-extracted features in the MPEG-7 format and export features from the database in MPEG-7 document format. To access/query

video data and experiment/test system functionality, we have developed a query interface that supports different kinds of queries, such as query by example image (matching image features against extracted video frame features or aggregate shot features) and query by motion (pan, tilt, rotate, zoom, etc.). We are building a web client as a web service to web-enable the entire system.

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

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

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

infrastructure, and has computation time within the time scale of network state updates of existing network services.

4. *Developing techniques for the large-scale storage of video data*. For physical storage management of multimedia documents, we have designed several novel data placement and scheduling schemes. These schemes are currently being implemented on a Sun E450 server and a Sun A1000 Raid array acquired through this grant. Also, due to the unique temporal nature of video data, we are analyzing several techniques for real-time disk scheduling. Managing large volumes of data necessitates the use of cheap tertiary storage. Given the very high random access cost of tertiary storage, efficient management of data is critical for performance. We are therefore developing data placement, migration, 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 Oracle9*i* Database Release 2. This development framework complies fully with the W3C XML data model, and absorbs the XML document model into the database for navigation and query. In addition to the native XML repository to the database, the Oracle XML DB encompasses both SQL and XML in a highly interoperable manner. With this combination, the application can store both structured and unstructured data that can be accessed using either SQL or XML operations interchangeably.

Networking

1. We have set forth a new unified theory of differentiated services that is implementable on IP networks. It improves on our earlier work on QoS scheduling, facilitating a theory of aggregate-flow QoS control.

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

3. We have installed the Purdue Infobahn comprised of four Cisco 7206 VXR routers that form an IP-over-SONET QoS backbone. We have implemented initial signaling and end-to-end controls using both RSVP (Int-Serv) and AS (Diff-Serv) to ascertain the routers' operating capabilities.

4. We have performed LAN-scale testing and benchmarking over the IP-over-SONET IP routers differentiated service provisioning, guaranteed service provisioning, and multiple time scale traffic control. Twelve PCs have dedicated connections to the testbed and are engaged as QoS flows and background traffic generators.

5. The present QoS testbed has been renamed to Q-Bahn (QoS Infobahn) to highlight its broad, comprehensive nature with respect to providing a complete solution for scalable QoS provisioning in IP internets. Q-Bahn consists of 9 Cisco 7206 VXR routers connected as an IP-over-SONET backbone, with 50+ PCs and workstations directly connected to the backbone for benchmarking. The physical topology of Q-Bahn is the same as that of Internet2/Abilene, albeit completely dedicated for advanced QoS

and security research, which Internet2/Abilene (due to its operational demands) is not. An even more distinctive feature of Q-Bahn is the operating system running inside the Cisco 7206 IP routers, called purdue-ios, which is a custom IOS (Cisco's router operating system) implementing the QoS switching algorithms developed in the Network Systems Lab. This unique collaborative effort between academia and industry led by Park (Directory of Network Systems Lab and co-PI of MSI project) allows advanced QoS research to be prototyped in a state-of-the-art production environment, a characteristic strength of the QoS component of the MSI project.

6. The Q-Bahn QoS testbed has been extended to incorporate a wireless mobile component, where a private six access point WLAN cellular network covering the three floors of the CS building is directly connected to the Q-Bahn testbed. The wireless mobile extension has been deemed necessary given the explosion, and expected dominance, of WLAN based local access technology, which injects additional complexities into provisioning effective end-to-end services that must be explicitly incorporated in the global system architecture. The wireless mobile access network is driven by a collection of in-house QoS-enabled applications developed in the Network Systems Lab, one of which, called QVI, a VoIP application incorporating QoS amplification mechanisms, runs over both Windows XP (PCs and laptops) and Windows CE (Compaq iPAQ Pocket PC). In addition to providing QoS-sensitive real-time VoIP and CD quality audio communication over integrated wired/wireless environments, the application QoS amplification infrastructure provides completely transparent QoS support for legacy applications via a QoS module, called Q-Driver, which resides in NDIS of the Windows XP and CE microkernels. Thus, our QoS platform is truly portable to environments characteristic of a typical end user.

7. We have generalized the unified differentiated services framework to a queuing framework (m-class G/G/1 queuing system), which will provide the most comprehensive theory of aggregate-flow scheduling yet and a significant advancement of scheduling theory.

Security

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

2. We have advanced and analyzed a denial-of-service (DoS) attack prevention framework based on probabilistic packet marking (PPM), and shown its effectiveness in the presence of single-source DoS attacks. We have used an adversarial framework to derive solutions to optimal decision making in a two-player environment comprised of the attacker and victim/target.

3. We have investigated the distributed DoS (DDos) attack problem under PPM and general network topologies where the attacker can pick the location and number of

attack hosts. We have shown that uncertainty factor amplification --- a measure of the attacker's ability to hide his true location(s) --- is achievable by reducing the attack volume at each individual attack site.

4. Our new work on scalable DDoS attack protection based on route-based distributed packet filtering (DPF) is being recognized as one of the very few effective and deployable DDoS protection mechanisms. One of our future goals is to incorporate the packet filtering mechanisms in the switching elements of the Q-Bahn testbed, thus providing further integration of QoS and network security. Two avenues we seek to explore are extension of the collaborative work with Cisco to include network security and exploration of Intel's IXP 1200 network processor as a programmable IP router platform wherein both QoS and security mechanisms can be implemented, tested, and benchmarked as part of the Q-Bahn testbed.

5. We have investigated the network security architecture of Cisco's router operating system IOS with the aim of incorporating the adaptive security architecture AdSec into its structure.