The findings of investigations into 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, are described below for the project period 2002-2003.

Multimedia Databases: Context-sensitive information retrieval, storage and dissemination

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.

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 WAN 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.

Security: 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.

Previous project periods

Multimedia Databases and Storage Management Systems

1. Content-Based Access to Video

Video Content Analysis. We have developed a novel video shot detection and key frame extraction technique that automatically adapts the threshold for detecting shots based on activities within video sequences. We have also developed a seed region aggregation and temporal tracking procedure to exploit video objects and their trajectories. We have designed and implemented a video preprocessing system with feature extraction functionality. The features currently extracted by our system are color histogram and camera motion. The HSV color space is used to extract the color histogram, and the dimensions of its components are: H (quantified with 18 levels), S

(quantified with four levels) and V (quantified with four levels). Nine types of camera motion are extracted and classified as camera motion features. Using this system, the video can be processed with shot segmentation, camera motion classification, hierarchical browsing, and video feature extraction. Based on the results of the video analysis, two categories of video retrieval can be performed with the system: (1) camera motion based video retrieval and (2) key frame based video retrieval. For motion based retrieval, users select a type of camera motion, and the system retrieves shots that fit the specified motion type. For key frame based retrieval, the user first browses the video collection to find an interesting frame to use for submitting an image-matching query. The user submits a query based on the frame and selected frame features. The system extracts the feature of the input frame online, and then uses the distance evaluation function to find all similar shots in the video database.

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

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

Metadata Presentation. Our system extracts new low-level features from the video that are specified in the MPEG-7 standard, including dominant color, scalable color, homogeneous texture, and edge histogram. Since significant development efforts have been directed towards the specification of the metadata (low-level and high level features) in a format that follows the MPEG-7 description scheme, we have developed a wrapper to import user-supplied MPEG-7 documents that are generated using multimedia description schemes for high level and low-level feature information. The document features are parsed and mapped to the system relational feature schema. The wrapper also exports feature information from the database as MPEG-7 documents. For the implementation of the wrapper, we use XML-DBMS and Java Packages for document transferral. We are also developing a JDBC driver for PREDATOR for database connectivity.

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

Querying Semi-Structured Data. We designed a constraint-based framework for querying semi-structured data. We provided a class of path constraints of interest in connection with both structured and semi-structured databases. Our constraint language is inspired by Feature Logics. Feature descriptions are used as the main data structure of so-called unification grammars, which is a popular family of declarative

formalisms for processing natural language. It provides a partial description of abstract objects by means of functional attributes called features. On top of the constraints we allow the definition of relations by means of definite clauses. The query language is based on the general scheme for handling clauses whose variables are constrained by an underlying constraint theory. Constraints can be seen as quantifier restrictions as they filter out the values that are assigned to the variables of a clause in any of the models of the constraint theory. The language is a hybrid one in the sense that it combines clauses with path constraints. It has a clear declarative and operational semantics.

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

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

2. Managing Video in Databases

Video Database (VDBMS) Prototype. We have introduced the real-time streaming of video as an integrated function within the database system. We are studying issues such as providing fixed bandwidth for each request and modifying buffer management to handle real-time as well as non real-time requests. We have identified some bottlenecks in current DBMS designs that are not suitable for processing video requests. We have also experimented on the current system under various workloads and conditions, such as large volumes of concurrent requests for different streams. We are investigating the scalability of our system to a growing number of stream requests. Another important aspect is the logging and recovery of video data. This requirement is a consequence of storing the raw video data inside the DBMS. We are investigating current limitations of logging and recovery with huge amounts of storage and long updates to binary contents. We have performed experiments to identify the effect of physical logging of video updates on overall system performance.

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

Query Processing and the Join Operator. We extract many high-dimensional features from the video data, on both the frame and shot levels. For each feature, we use a high-dimensional index structure as the access path for similarity queries on this feature. To enable the database management system for querying videos based on multiple features and querying by multiple examples, the similarity results from each index

structure need to be joined in a way that gives a global similarity ranking over all specified features. We have enhanced the query processing capability of the system by introducing a new rank-join query operator to handle these queries. The new join operator is a binary, pipelined operator that advances the optimal aggregation algorithms introduced in the literature by transforming an optimal theoretical solution into one that can be implemented in a practical database engine.

Stream Query Processing. We define a stream as an infinite sequence of data items, where items are appended to the sequence over time. Videos can be considered streams of large objects (video frames and audio-related data). We address the issue of executing online queries on multiple video streams for the purposes of video stream editing, analysis, and monitoring. As an example, we might need to hide the faces of patients in a medical education video of patient consultations in order to protect patient privacy. We address the online processing of video since we may not want to store the preprocessed video, and we do not want to create a special edited version of the video for each patient. As a second example, consider the tracking of persons between two video cameras. Online video preprocessing is used to extract "face" objects that appear in each video stream and determine their identify 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 manger and the query execution plan. Initially, the StreamScan contacts the Sensor Table and registers stream requests to the stream manager. This triggers the stream manager to start retrieving sensor values. Based on a request from other query operators, StreamScan retrieves a tuple-at-a-time from the buffers of the stream manager for processing by the query execution plan. We have already implemented two novel join techniques for online sensor data in the query engine: the nested-loop window-join and hash window-join. We are currently investigating the issues of (1) query execution with limited amounts of memory, (2) intermediate storage, and (3) scalability (queries with large number of streams). We believe that streams of large objects have additional complexity that imposes further challenges on stream query processing.

Online Video Processing. To address the online processing of video, we have implemented video processing functions such as blurring and fast forwarding. Those functions are executed within the query plan and work on the video in the compressed

domain so that it is not necessary to decompress, process, and then re-compress the video to produce the final online results.

3. Quality of Service in Video Databases

Quality-of-Service Aware Repository (QuaSAR) Architecture. We have designed a system architecture that supports user quality-sensitive queries within a database framework. The proposed architecture relies upon the notion of QoS aware interfaces to the various components of the system, such as the network layer and the operating system (encompassing CPU, main memory, and disk storage). These interfaces enable real-time determination of the status of the components with respect to the satisfaction of QoS constraints. In addition, these interfaces will support reservation of resources to guarantee the ability to satisfy the user's requested level of quality. A key component of QuaSAR is the enhanced query processing capabilities in contrast to traditional databases. Based on the content component of the query and the content metadata, alternative plans are generated for the retrieval of the relevant objects. Each plan is annotated with QoS parameters relevant to each component based upon translation of the user's quality parameters for the given plan. Each of the constraints represented by the annotations are tested through the interfaces and, if necessary, reserve resources. If no feasible plan is found, a negotiation step is invoked to adjust the constraints and re-evaluate the feasibility of the plan.

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

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

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

4. Large Scale Video Storage

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

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

Networking

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

Configure Purdue Infobahn QoS Testbed and Perform Iinitial Tests and Exploratory Iimplementations. As a starting point to implementing a full-fledged version of our

optimal aggregate-flow classifier inside Cisco 7206 VXR routers with system support from Cisco, we have implemented end-to-end signaling controls, both for current IETF Diff-Serv specifications (i.e., Assured Service) and IETF integrated services (Int-Serv) using RSVP. Our end-to-end signaling platform can be adapted, with modular changes, to the optimal aggregate-flow per-hop control architecture advanced in our theoretical and simulation work. We have tested the real-time scheduling and switching capabilities of a network of Cisco 7206 VXR routers connected as an IP-over-SONET internetwork with 100 Mbps network interfaces tapping end station PCs—serving both as QoS end points and traffic generators—and ascertained their operating dimensions and properties with respect to yielding predictable performance. The latter include saturation effects stemming from processor scheduling and software switching overhead at the router. We are in the process of implementing the first version of optimal aggregate-flow PHB inside Cisco 7206 VXR routers, which will then be tested with respect to its QoS provisioning properties when incorporating all software and processing overheads associated with routers and end stations.

Advance End-to-end Traffic Control for QoS Shaping. We have developed and implemented end-to-end QoS amplification techniques using adaptive packet-level forward error correction and multiple time scale traffic control. Adaptive FEC (AFEC) allows invariant end-to-end QoS to be exported over variable network conditions stemming from shared bottleneck routers and congestion effects. Since differentiated services-even in the presence of optimal aggregate-flow PHB-is subject to occasional variability due to imperfect QoS protection resulting from abstaining from per-flow reservation and admission control, AFEC can provide QoS amplification over varying degrees of imperfect QoS channels which can further improve the operating range and usefulness of differentiated services on top of scalability. We have advanced the multiple time scale (MTS) traffic control framework that exploits large time scale correlation structure present in bursty-in particular, self-similar-network traffic for traffic control purposes. An important consequence of MTS traffic control is the ability to mitigate the cost of feedback traffic controls in high delay-bandwidth product networks that is especially pronounced in high-bandwidth wide area networks. By exploiting predictability structure resident at time scales exceeding the round-trip time (RTT) or feedback loop by an order of magnitude or more, timeliness of reactive can be achieved by bridging the 'uncertainty gap' inherent in reactive actions. Thus, MTS traffic control facilitates a much needed measure of proactivity, which, in turn, facilitates scalable end-to-end traffic and QoS control. MTS traffic control has been implemented both for QoS (a multiple time scale extension of AFEC called AFEC-MT), and throughput maximization for TCP (MTS extension called TCP-MT over Reno, Tahoe, and Vegas) and rate-based control for ATM. Multiple time scale redundancy control (i.e., AFEC-MT) has been benchmarked over UDP/IP internetworks that transport real-time MPEG I & II video and audio. The systems run both for UNIX and Windows NT. In the latter, the sender interfaces with a real-time MPEG compression board (Optibase I for MPEG I video/audio and Futuretel for MPEG II video/audio), which is fed by high-end digital and analog video cameras.

We have implemented and tested both differentiated services and guaranteed services functionalities on Purdue Infobahn. We have observed that RSVP-based per-flow resource reservation shields user flows from the detrimental effects of cross-traffic whereas for differentiated services the achieved service level is more variable. The influence of processing speed of the router can be significant when subject to high load, and it is a characteristic that needs to be explicitly incorporated when computing end-to-

end QoS assurances. The first service parallels the standardization effort under the auspices of the Internet2 QoS Working Group (co-PI Kihong Park is a member of the Architecture Design Team) which seeks to achieve robust Premium Service using IETF's Expedited Forwarding (EF) specification, however, using conventional per-flow reservation.

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

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

Security

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

We have advanced a framework for probabilistic packet marking for distributed denial-of-service attack (DoS) prevention that puts forth an effective technique for dealing with an

important security and QoS threat. Instead of logging path information at routers which makes traceback of DoS attack overhead-prone and ineffective in high-speed networks where routers are expected to switch at Gbps or higher rates, the complexity is pushed to the edge following the end-to-end paradigm, and 'logging' is carried out by probabilistic sampling of path information in an end-to-end manner. This method can be easily implemented in IP routers. Our work provides a comprehensive analysis of the power of probabilistic packet marking (PPM) with respect to DoS, both single-source and distributed. We use a 2-player adversarial framework to analyze the minimax and maximin optimal strategies of the attacker and victim, and show that PPM is effective at preventing DoS attacks. We show using numerical results the outcome of optimal strategies when using typical Internet based configurations including hop count, topology, spoofing, and marking probability.

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