1 Research Description

1.1 Overview

Motivation The past few years have seen a tremendous proliferation of broadband networks. These networks have the potential of providing opportunities for attractive multimedia applications, including video-conferencing, distance learning, digital libraries, entertainment, commerce, command and control systems, tele-medicine etc. Supporting the diverse quality of service (QoS) requirements of such applications is a critical and challenging problem involving database, networking, and data security technologies. QoS constraints of applications may be expressed in terms of timing, storage, processing, security, and presentation of multimedia data. Various components of the multimedia infrastructure must be capable of satisfying these constraints individually, as well as collectively for end-to-end performance. The database component must have capabilities to store, process, and retrieve information distributed across networked environments. This requires QoS-sensitive network services that facilitate timely distribution and collection of distributed multimedia information. Translation of user requests with quality of presentation (QoP) constraints to internal representations and actions that are executed by individual components is another nontrivial problem. For example, in applications that involve the transfer of video with timing and quality constraints, disk bandwidth at source, network bandwidth at routers, and CPU cycles at end stations need to be scheduled and matched to achieve target QoS. This involves careful negotiation between QoP and QoS parameters with respect to the overall system resources.

Large-scale multimedia systems may be required to support thousands of concurrent accesses to QoS sensitive data. This may cause crippling network congestion even in Gbps backbones. Under such circumstances, protecting users and applications with guaranteed services from congestion effects and provisioning soft QoP to applications with rigid QoS network constraints is a challenging problem. Data security in the presence of distributed servers is essential for real-world applications. Sensitive multimedia data in such applications must be protected by encryption and supporting mechanisms for access control, authentication, and watermarking. These services often incur significant overheads and must be explicitly incorporated into QoP-based multimedia data scheduling. The proposed research will attempt to provide efficient and robust solutions for the crucial problems mentioned above.

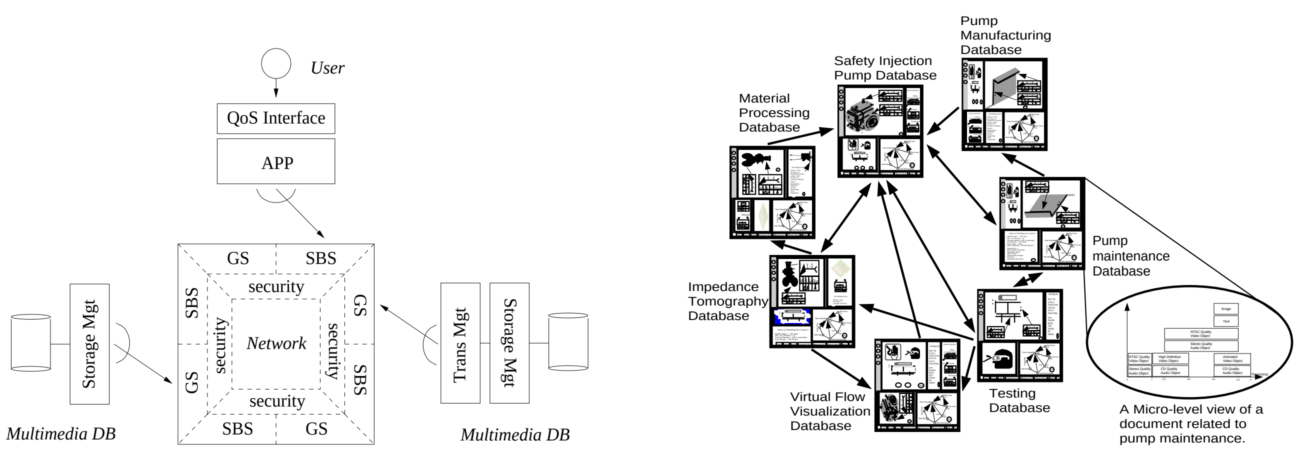

Figure 5: (a) Overview of proposed MSI Architecture. (b) An example two-level, for multimedia nuclear reactor safety database.

New Challenges This proposal addresses the issues of provision of desired levels of QoS and QoP for distributed stored, networked, and live multimedia data. The proposed hardware and software novelties will allow seamless integration of QoS management across the three principal areas – databases, networks, and security. Each of these areas presents significant research challenges. The five PIs–Elmagarmid (databases), Ghafoor (multimedia), Korb (multimedia testbed construction), Park (networks), and Spafford (data security)– along with their colleagues are actively engaged in cutting-edge research in these areas with significant funding from NSF and other agencies. The proposal also outlines extensive ongoing and planned collaborative research efforts that are critical for addressing the problem of integrated QoS management. The proposed research infrastructure comprises of state-of-the-art equipment in networking, storage, capture, compression, presentation and development. This will allow us to build a campus-wide testbed for integrated QoS management research for distributed multimedia databases. The infrastructure will be used for implementing, testing, and benchmarking new architectures, algorithms, interfaces, protocols, translation mechanisms, and applications for QoS-based scheduling . We expect the synergy of our coordinated research activities under the unifying umbrella of quality of service to bring forth significant new advances in the current state-of-the-art. These advances will have wide-ranging impact on multimedia services, with implications for Internet2 and the facilitation of an efficient information society at large.

Organization This proposal is composed of three parts – multimedia databases, networks, and security. Within each part, we describe on-going research along with proposed collaborative and individual research. The proposed MSI will be used to build advanced architecture that will meet the above-mentioned challenges. Each part concludes with a set of proposed experiments to be carried out over the Purdue-wide research testbed and subsequent expansion into nation-wide broadband environments such as vBNS and Abilene where Purdue is a member institution.

The initial applications that will drive the MSI research testbed will include: real-time video-conferencing, video-on-demand, scalable data compression, and distributed multimedia database management systems. These will form building blocks for a number of applications such as distance learning, digital libraries, real-time monitoring, and information broadcast systems. We plan to use multimedia data from real-life applications, including nuclear reactor safety research and veterinary medicine. These applications are motivated by their massive multimedia data requirements. We propose to build a digital library and live video source exporting QoS-sensitive access and transmission services customized to their needs. This will also allow us to ascertain the flexibility and configurability of the MSI in terms of QoS/QoP-sensitive data access and delivery.

2 Distributed Multimedia Databases

We envision that emerging multimedia applications will require the support of distributed multimedia databases connected over wide-area broadband networks. Such systems will support distributed browsing environments [13], where users can browse through complex multimedia objects consisting of text, video, audio, image etc., retrieved from distributed servers. The overall information space can be partitioned into several classes of multimedia objects and linked together in form of a two-level browsing graph. Figure 5(b) shows an example of such a multimedia browsing environment for nuclear reactor safety applications. In this figure, the first level provides directed navigational links among several complex multimedia objects. The second level represents the object being accessed along with the detailed information about the objects' meta-schema consisting of several attributes. These include spatio-temporal structure, a set of quality of presentation (QoP) parameters, media-specific processing functions, access mechanisms etc. A typical multimedia application may comprise of a large number of such browsing graphs maintained in the form of a global multimedia database meta-schema. Such browsing graphs are created by the users for possible classification of complex multimedia objects with relevant topics and contents.

Currently we are pursuing two basic areas of research related to this system. These are end-system database server architecture, and distributed multimedia data management. We highlight leading research tasks and describe the benefits of experimentation related to our ongoing research using the proposed MSI.

2.1 End-System Database Architecture

The major role of a multimedia server is to perform the critical database and operating system functions to ensure the highest levels of quality of service for multimedia applications [33]. There are four major functionalities that are inter-linked and affect the performance and the quality of service provided by the multimedia database management system. These include the physical storage management in terms of data placement, compression of multimedia data to conserve system resources, I/O management and disk scheduling techniques to achieve high-performance, and a mechanism to facilitate for the users the specification of high-level quality of presentation parameters. The users specify their desired QoP in terms of the database meta-schema. The meta-schema along with the compression techniques establishes resource requirements (in particular, I/O bandwidth), allows the development of intelligent disk scheduling algorithms, and forms a basis for the design of efficient data placement strategies. The ongoing research by the investigators in this area and its relevance to the MSI is described below.

2.1.1 Rate-Scalable Video Compression

Many applications require that digital video be delivered over computer networks. Network data bandwidth limitation, however, almost always poses a problem [30]. A user may request a video sequence with a specific quality. One solution to this problem is to compress and store a video sequence at different data rates. The server then delivers the requested video at the proper rate given network loading and the specific user request. This approach requires more resources to be used on the server in terms of disk space and management overhead. Therefore scalability, the capability of decoding a compressed sequence at different data rates, has become a very important issue in video coding.

A specific coding strategy known as embedded rate scalable coding is well suited for continuous rate scalable applications. In embedded coding, all the compressed data is embedded in a single bit stream and can be decoded at different data rates. We have developed a new continuous rate scalable hybrid video compression algorithm using the wavelet transform, called SAMCoW. It allows the decoding data rate to be dynamically changed to meet network loading. We intend to deploy SAMCoW on MSI and compare its performance on the network with MPEG-2. In particular, we want to examine whether a true rate scalable video compression technique performs better in a network with real QoS constraints when compared with off-the-shelf techniques such as MPEG-2.

We intend also to explore how the new MPEG-4 and MPEG-7 standards can be used in our system [11]. MPEG-4 addresses the issue of how one codes or represents multimedia contents. The key technology breakthrough of MPEG-4 that makes it fundamentally different from MPEG-2 has been the ability to encode visual objects of arbitrary shape. The question that has not been addressed is how this new, and very different, approach to multimedia data compression can be used in the context of a networked multimedia system. An issue that need to be addressed is can content scalability be exploited relative to QoS demands from a user and/or the network. MPEG-7 will specify a standard set of descriptors that can be used to describe various types of multimedia information. We feel that our proposed system can act as a testbed for investigating how the concepts developed in MPEG-7 can be used in a real system. For example, how can the MPEG-7 descriptors be used a loaded network environment? This research is being carried out by Bowman, Delp, and Grama.

2.1.2 Meta-Schema Design for Multimedia Objects

We have developed two working prototypes, one of a complete video database management system [13] and the other of a similarity based browsing environment [4]. These two systems operate independently and one of our goals will be to integrate these two prototypes as they have complementary functionalities [9]. Though, the goal of these two prototypes has been to store, index, retrieve, and query video data based on semantic content, we have developed a conceptual data model that is particularly interesting for the type of applications being addressed here. This novel data model, called VideoText, is a multi-modal video data model [14, 12]. This model is defined over video streams, various annotations along the different media modalities, and maps that relate these annotations to particular segments of video streams. This characteristic will prove very useful to the two level browsing environment we will develop in this new MSI. We will also proceed to develop new query languages, and new compact representations of video streams. We will also use scene change detection and identification in order to segment video into logical segments. A process of shot descriptions using pseudo-semantic classes defined by the database designer will then be used to classify the shots. Finally we use the active browsing facility based on similarity pyramids to access and filter the data.

We plan to compare the efficacy of our own semantic based data model with existing conceptual models. These other models can be classified into several categories such as language based models (e.g. HyTime), graphical models, Petri-net based models, and temporal abstraction models. Some models are primarily aimed at synchronization aspects of multimedia data, while others are more concerned with the browsing aspects of multimedia objects.

We are currently analyzing these models in terms of their capabilities to represent key attributes of the database meta-schema. These attributes include: spatio-temporal structure, quality of presentation (QoP) parameters, media-specific processing functions, and object retrieval mechanisms. The spatio-temporal structure describes the space/time synchronization requirements for complex multimedia objects [21]. One leading model is based on Petri-nets and incorporates the QoP parameters to specify the desired quality of multimedia presentation which include reliability, level of resolution for different types of multimedia data, rate of presentation of isochronous data etc [1]. The media specific functions are used to carry out user's specified processing prior to the presentation. These functions may be needed for data conversion, enhancement of displayed data, coloring of images or video, dubbing etc. The object retrieval mechanisms facilitate the storage sub-system in terms of disk scheduling to satisfy real-time object delivery requirements.

The other key issue in our analysis of these models is to compare their suitability for indexing complex multimedia objects. For example, in a language-based model such as HyTime, indexing and querying of objects is clumsy. In this case, the language description needs to be transformed into appropriate data structures, and this requires sophisticated parsing and transformation tools. The research related to this topic is being carried out by Bowman, Delp, Elmagarmid, and Ghafoor.

2.1.3 Intelligent Multimedia Data Placement Strategies

Physical data placement is a critical issue in the performance of interactive multimedia information systems. We propose to develop a systematic approach to automatically generate storage layouts for hierarchical storage (consisting of primary, secondary and tertiary storage) in order to support a complex multimedia object browsing environment. For the two-level browsing space, we plan to model the access pattern of navigation as a time-homogeneous ergodic Markov Chain. In essence, this model captures the behavior of user interactions such as cyclic browsing, access statistics, and temporal relations among multimedia objects. From this model, the stationary probability for each node (representing a complex multimedia object) of the navigation graph can be found. Subsequently, a large multimedia data placement problem can be divided into a number of small problems by decomposing the navigation graph using a partitioning algorithm, such as Fiduccia-Mattheyses approach [10]. The partitioned graph can be distributed over multiple disks by assigning the objects in the same partition to a single disk or a small group of disks. Our initial analysis proves that the problem of multimedia data layout is an NP-complete problem, for which no polynomial time solutions are known. For this problem, we have developed and analyzed several heuristics to solve this problem using stationary probabilities and expected access costs [7]. Our simulation study has revealed that the maximum cost chain connection algorithm can generate a layout close to optimal placement for browsing graphs of the type shown in Figure 5(b). This is particularly true for browsing graphs with low connectivities. The research related to this topic is being carried out by Ghafoor and Prabhakar.

2.1.4 High-Performance I/O and Disk Scheduling Algorithms

I/O and disk scheduling has traditionally been a limiting factor in the performance of conventional database servers. The I/O bottleneck becomes even more pronounced for currently evolving systems handling multimedia data, such as audio and video data, for which the I/O bandwidth requirements can change drastically over time. Since I/O bandwidth is a precious resource, we plan to design and evaluate efficient techniques for I/O management and synchronous retrieval of such data. Such schemes also provide an admission control mechanism for allowing new sessions. The QoP parameters associated with complex multimedia objects will be the key consideration in designing such techniques.

A basic QoP requirement is to ensure inter-media and intra-media temporal synchronization at the user-interaction level. We are currently we exploring several I/O and disk scheduling schemes that use the notion of QoP. Such schemes can be broadly divided into two categories: static and dynamic. In static schemes storage sub-system resources, especially the I/O channels, are statically allocated to support retrieval of concurrent multimedia streams. Allocation of resources is treated as an optimization problem. The static resource allocation mechanisms require intelligent disk scheduling algorithms, whereby access and retrieval of isochronous data needs to be carried in real-time. We have recently shown that the problem of disk scheduling for multimedia data retrieval system to be NP-hard [16]. We have designed several deadline driven heuristics to solve this problem. Under resource constraints, such as the I/O bandwidth or the size of the main memory, the disk system may need to filter data, whereby high-resolution components of video and image data may be skipped. In addition, we are developing several heuristics for disk scheduling that exploit intelligent data filtering schemes. The objective is to study the tradeoffs between I/O bandwidth utilization and the desired QoP.

In dynamic schemes, resources such as I/O bandwidth are demand-driven. The major objective of such schemes is to maximize the resource utilization such as the I/O bandwidth. In these schemes, QoP parameters associated with multimedia objects are used in order to estimate dynamic I/O channel requirements and the size of the main memory needed to support concurrent sessions. The general approach is to model the I/O channel allocation as an optimization problem where the I/O channel is a scarce resource. Under such conditions, the disk system can perform essential data filtering functions, within the bounds dictated by the QoP parameters. We are designing several schemes and analyzing their performance in terms of their complexity, and the achievable QoP. The research related to this topic is being carried out by Ghafoor and Prabhakar.

2.1.5 Fault Tolerance and Scheduling of Tertiary Storage

Due to the extremely high latency of tertiary storage, it is necessary to manage tertiary I/O judiciously. Performance enhancements can be achieved through I/O scheduling and careful data migration between disk and tape. Despite the importance of I/O scheduling, only schemes which handle one or two media have been developed. We have developed a near-optimal scheduling algorithm for robotic libraries [29]. Although this algorithm achieves significant improvement over existing schemes, it is not tailored for multimedia data. We will deploy the existing algorithm and investigate better algorithms for multimedia data. Typical migration of data from tape to disk is demand-driven and can result in large latency delays. We will develop anticipatory migration schemes which pre-fetch data from tape based upon user browsing patterns (using hints from the browsing graph). Fault tolerance is typically handled at the secondary storage level. However, when the data resides primarily on tertiary storage, fault tolerance must be supported at the tertiary level. Standard techniques such as mirroring and RAIT (analogous to disk RAID) do not work well on tertiary storage due to their high storage overhead or the need for synchronization between drives. We have developed a novel scheme for avoiding both these problems by storing the data and parity asynchronously [15]. We will deploy this new scheme and study how it impacts the need for fault tolerance at the secondary storage level. This research related to this topic is being carried out by Prabhakar.

2.2 Distributed Multimedia Databases

Emerging multimedia applications will require networked multimedia databases. An important consideration in this environment is to provide seamless integration of these databases and underlying networking infrastructure to support QoS-based multimedia applications. Such integration poses several technical challenges [9]. These include the following [2, 19]:

- quality of service management for distributed multimedia database applications. In particular, mapping of application level QoP parameters to network QoS parameters entails determination of resource allocation and admission control policies for the overall system;
- designing synchronization mechanisms to compose complex multimedia objects from distributed servers. Distributed data caching is expected to play an important role to facilitate synchronized delivery of complex multimedia objects;
- location identification and global access mechanisms for integration of information.

We have been pursuing these research issues and have proposed several techniques to solve these problems. In particular, in [20], we have proposed a hierarchical multimedia composition mechanism using the concepts from distributed systems and networking. We are currently researching the QoS-QoP negotiation protocols aimed to find the "best" compromise between capabilities and constraints of the database servers, the network and the clients, and particular user requirements and preferences, expressed as part of QoP. We are pursuing development of distributed object access mechanisms with particular consideration given to effective methods for browsing and querying large distributed multimedia databases.

Integration of distributed multimedia databases involves multiple sites, of multiple data types, or on multiple platforms. The distributed database servers are expected to be primarily object-oriented, supported by stored procedures for robustness and scalability. This research related to this topic is being carried out by Ghafoor and Elmagarmid.

2.3 Proposed Research and Experiments

For experimental prototyping of the distributed multimedia database system, we are planning to use data from the nuclear reactor safety database depicted in Figure 5(b) and the veterinary medicine video data. There will be extensive experimentation of the prototype in terms of further pursuance of the above mentioned tasks. The details of such prototyping and perceived benefits are given below.

2.3.1 Meta-Schema Development and Benchmarking of Multimedia Database Servers

For the video database management and the similarity-based browsing environment discussed in the previous section, we have developed a conceptual data model. We plan to evaluate this model and compare its performance with respect to other models in terms of QoP expressiveness, query processing support, and building active browsing facilities. In order to get meaningful performance results, we plan to test and benchmark several such models for the selected applications related to nuclear reactor safety. Unlike traditional database applications, benchmarking for multimedia applications is still an open research issue. The imprecise nature of multimedia data and queries, makes it difficult to define correctness criteria for retrieved data. Moreover, the correctness of the results of a query is context dependent – different applications or users may not have the same definition of correctness. An important criterion for evaluating multimedia applications is therefore the quality of the result set, in addition to the traditional speed or efficiency measures.

2.3.2 Implementation of Data Placement Strategies for Hierarchical Data Storage Systems and Disk Scheduling Policies

Using the proposed MSI, we plan to implement and evaluate the data placement strategies discussed in the previous section. The schemes will be implemented on a hierarchical storage system consisting of magnetic disks at the secondary level and a robotic library at the tertiary level. Data has to be migrated from tertiary to secondary storage before it is delivered to the network. Each such migration presents an opportunity to generate a new placement of the objects on secondary storage taking into account the latest information about access patterns. We will evaluate the performance of these placement schemes against more traditional approaches, including a standard RAID scheme and a commercial multimedia data placement schemes. The knowledge of the browsing patterns will also be used to guide the placement of data on tertiary storage. For example, objects that are clustered together based upon the graph partitioning may be stored on a single tape to minimize switching. In our experiments we will evaluate several multimedia data placement schemes for secondary and tertiary storage. The results expected from these experiments include a set of data placement schemes for browsing complex multimedia data on hierarchical storage systems.

In particular, we will investigate the following three alternatives for implementing the secondary storage layer. 1) An off-the-shelf RAID system since it provides high performance at low cost. 2) A collection of disks under the control of a commercial storage manager and scheduler that explicitly handles continuous media [23]. 3) A collection of disks under the control of the storage manager that will used to implement our disk scheduling schemes. The rationale behind these choices is that the RAID system is a typical solution chosen for secondary storage management. However, we believe that performance enhancements can be achieved by tailoring disk scheduling and management strategies to specific application domains. The CineBlitz storage manager is an example of this approach for the multimedia domain. We will employ CineBlitz for implementing option 2. CineBlitz does not, however, address the issues of data stored on tertiary storage. We plan to employ the tools and techniques developed by us previously in order to investigate option 3. In this manner, we will compare the three options.

2.3.3 Distributed Multimedia Databases

For the proposed MSI, we plan to interconnect several multimedia databases over the broadband network. We have been pursuing research in the area of distributed multimedia databases and have proposed several solutions to tackle important technical issues in this area. We plan to perform several experiments on the proposed MSI in order to to analyze and validate our research results. In particular, we plan to conduct extensive experiments for: (i) implementing mapping mechanism between application level QoP parameters to network QoS parameters and interfacing with the the underlying networking sub-system in order to evaluate various resource allocation and admission control policies for the overall system; (ii) developing a distributed directory management system for locating and accessing complex multimedia objects, and (iii) developing and evaluating hierarchical multimedia synchronization mechanisms using distributed caching.

3 Networking

The following sections give a summary overview of the Integrated Guaranteed Service and Stratified Best-Effort Service Architecture and a list of proposed research and experiments to be conducted over the Research Infrastructure testbed.

The QoS architecture research is supported, in part, by two NSF grants—ANI-9714707 (Park) and ESS-9806741 (Park, Yau)—and grants from the Purdue Research Foundation and Sprint (Park). Related networking research funded by NSF and other sources that will benefit from the proposed Purdue-wide implementation of the integrated QoS architecture are summarized in the Proposed Research section. The integrated QoS architecture research is conducted across two labs at Purdue University—the Network Systems Lab (Park) and the Systems Software and Architecture Lab (Yau). The summary overview of the Integrated Guaranteed Service and Stratified Best-Effort Service Architecture follows [5, 6, 27, 26].

3.1 Motivation for Dual QoS Network Architecture

Quality of service (QoS) provision in high-speed networks carrying a diverse mixture of traffic from e-mail to bulk data to voice, audio, and video is a difficult problem. The traditional approach is to reserve resources along a route to be followed by a traffic stream such that the latter's burstiness can be suitably accommodated. One may classify applications as reserved or best-effort based on whether their QoS requirements are "stringent" or not.

• Reserved Traffic The dividing line between guaranteed and best-effort service lies in the current limitation of control algorithms to achieve stringent quality of service guarantees without resorting to conservative resource reservation—i.e., over-provisioning. Although research abounds [8, 24] (we omit further references due to space constraints), analytical tools for computing and provisioning QoS guarantees rely on resource reservation coupled with careful admission control which, in turn, depend on input traffic shaping to preserve well-behavedness properties across switches that implement a form of generalized processor sharing (GPS) or weighted fair queueing packet scheduling.

The scale-invariant burstiness associated with self-similar network traffic [17, 28] limits the shapability of input traffic while at the same time reserving bandwidth that is significantly smaller than the peak transmission rate. The performance impact of self-similar traffic—which, unlike Poisson sources, does not smooth out with aggregation—has been shown to be detrimental and this imposes a trade-off relationship between QoS and resource utilization which limits the degree of utilization achievable while guaranteeing stringent QoS. This, in turn, renders guaranteed services prohibitive with respect to their provisioning cost.

Another important factor is scalability. For guaranteed traffic, resource reservation and admission control is administered on a per-flow basis and this entails, among other things, that per-flow state be maintained at every router in the system. As the size of the system—i.e., number of users, hosts, and routers in the system—increases, maintaining per-flow state information and computing feasible resource allocation using reservation and admission control for thousands of users becomes a difficult problem, increasing both the complexity and processing overhead associated with network elements.

• Stratified Best-Effort Traffic Given that provisioning guaranteed service is a necessary service for highly stringent QoS-sensitive applications such as life critical remote monitoring (e.g., surgery assisted by experts receiving real-time video/audio feeds at remote sites) but provisioning day-to-day QoS using reservation and admission control incurs a prohibitive cost for the average user, it is necessary to consider imposing a service class structure on best-effort bandwidth such that a diverse spectrum of *graded* services can be exported to noncritical yet still QoS-sensitive applications. It would be overkill and wasteful to provision such QoS needs over reserved channels. On the other hand, relying on homogenous best-effort service, characteristic of today's Internet, would be equally unsatisfactory. When reserved and nonreserved traffic coexist, the latter may also help amortize inefficiencies stemming from over-provisioned resources for guaranteed traffic and thus increase system utilization and reduce the cost of QoS-sensitive services, even those of guaranteed services.

3.2 QoS-Sensitive Services Exported by the Network

The network exports two types of services—guaranteed and stratified best-effort—to the end user and DB component. Security services, in our architecture, are *imported* by the network in the sense that the user or DB's security requirements are part of a *generalized QoS requirement vector* specified by the user or DB component, and are implemented, in part, by security service modules in the network protocol stack.

The user's QoS requirements are conveyed to the network using a generalized QoS requirement vector—or simply QoS requirement vector (QRV)—whose components include bounds on end-to-end delay, jitter, packet loss rate, throughput, encryption option, authentication option, reliability measure, and a host of other indicators. The components of a QRV are classified into three types—QoS, security, and reliability—where the latter represents the degree of stringency with which a particular QoS level (e.g., delay bound) is to be met.

• Guaranteed Service Guaranteed service (GS) is accessed by users or DB components via the GS service access point (GS-SAP). The client, along with the service request, passes along a QRV which captures the user's desired target QoS. The *semantics* exported by GS are as follows. Given the user or DB's service request and its associated QRV, if the network accepts the request, it *guarantees* that all the requirements specified in the QoS contract will be met in the sense that unless events beyond the network's control such as earthquakes, fires, and other calamities occur which alleviate the network service provider's responsibilities analogous to insurance contracts, the network's mechanisms will execute the service request commensurate with the desired QoS or better. In general, the guarantee is met with a certain probability—i.e., reliability—assuming it can be computed.

• Stratified Best-Effort Service Stratified best-effort service (SBS) provides QoS-sensitive services as in GS—albeit without the latter's strict guarantees—and is accessed via a well-defined SBS service access point (SBS-SAP). The client submits a service request along with a QRV which contains the same set of QoS indicators as in GS. However, unlike in the latter, there is no admission control and resource reservation in SBS. The semantics exported by SBS do not provide guarantees in the sense of GS—i.e., absolute protection modulo a set of pre-defined exceptions—rather, it provides a weaker form of service contract which *induces* protection from the adverse effects of resource contention but does not guarantee complete impunity. The service contract provided by the network to the user can be represented by the following three behavioral rules: (i) *selfishness emulation*, (ii) *least cost service*, and (iii) *zero-sum fairness*. We omit their explanation due to space constraints [5, 6].

3.3 Guaranteed Service Architecture

We provide guaranteed services using resource reservation and admission control. We follow the approach of [8, 24] for virtual circuit switched networks which uses resource reservation and admission control coupled with leaky bucket traffic shaping/policing to deliver deterministic guarantees on QoS. The main building block of guaranteed services are GPS routers that enable the reservation of resources at a switch on a per-flow basis. The principal functional components of GS are as follows:

• RSVP-Compliant GPS Routers We will use RSVP as the underlying protocol for reserving resources at GPS routers along a path or multicast tree. The equipment budget provides for six 72xx-level Cisco routers for this purpose. They will constitute the primary Purdue-wide backbone network for QoS-sensitive packet switching. The ten FORE ATM switches will be used as departmental aggregation points and either tunneled through via AAL 1 and 2 or interfaced with FORE's Intelligent Infrastructure layer to achieve guaranteed QoS.

• End Stations with Real-Time OS A necessary component to providing end-to-end QoS guarantees is real-time CPU scheduling capability at end stations. We will use a kernel-based modification of Solaris UNIX which uses the ARC real-time CPU scheduler to provide guaranteed CPU cycles at end stations. An RSVP capable prototype for UltraSparc and Pentium workstations has been developed at the Systems Software and Architecture Lab (Yau) and is operational.

• Admission Control We will use a form of reservation-based admission control using service curves to compute deterministic guarantees for bursty input traffic using leaky bucket traffic shaping/policing. We will use QoS-sensitive routing to compute minimum delay/maximum bandwidth paths and interface with Cisco IOS Layer 3 signalling protocols to control VC connection set-up and management. A prototype system has been developed at the Network Systems Lab (Park) and is being ported to run on the DUNES distributed operating system.

3.4 Stratified Best-Effort Service Architecture

The stratified best-effort service (SBS) architecture is built using the same building block as GS—GPS routers—and as such complements, and is compatible with, the latter. A full-fledged system for WANs (e.g., vBNS) has been built and tested using large-scale simulations at the Network Systems Lab (Park). SBS is implemented using the SBSP protocol which is installed at all participating GPS routers in the network. SBSP is also run at end stations on a per-flow basis. Inside the network, however, SBSP is completely per-flow state independent and performs a simple service class index computation as part of distributed QoS control. The principal functional components of the SBS architecture are as follows:

• Distributed QoS Manager The SBSP has two components, one of which is the distributed QoS manager (DQM) which is a distributed QoS control module that runs at every GPS router in the system. Mechanistically, all that a DQM does is intercept incoming packets to a router, inspect the flow's end-to-end QoS requirement, decide what service class to assign the packet at the *local* switch to achieve a desired target QoS—i.e., local QoS responsibility—and forward the packet with a newly enscribed service class label to the GPS switch proper at which it is serviced according to its service class label. Thus, as far as the GPS switch concerned, SBSP is completely transparent.

• End-to-End QoS Manager The second component of SBSP—the end-to-end QoS manager—has two responsibilities, one, to monitor the rendered or measured end-to-end QoS and feed it back to the sender for end-to-end feedback QoS control, and two, provide an access point through which the flow's internal resource usage is tallied. The latter is achieved at a fine-granular level with the help of a service class-QoS association table maintained at every router which monitors the measured QoS of every service class—a task that most routers perform anyway—and also assigns prices to service classes according to their rendered QoS: the higher the QoS of a service class the higher its unit cost. This information can then be accessed by a service provider to facilitate usage pricing and other forms of encapsulation.

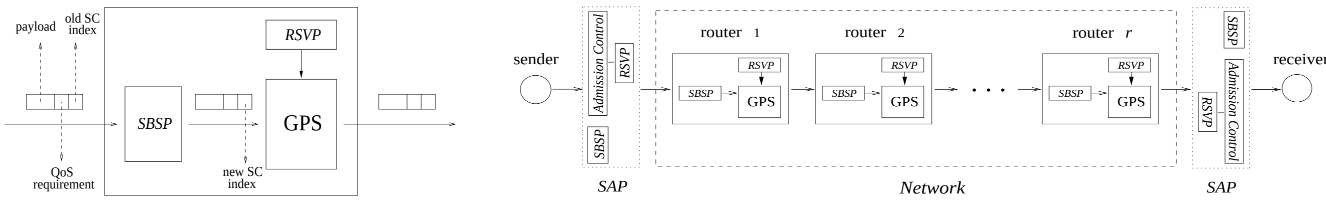

Figure 6: Left: Integrated RSVP/SBSP-augmented GPS router. Right: Integrated GS/SBS service provisioning using RSVP/SBSP-augmented GPS routers.

Figure 6 (left) shows a diagram of the integrated RSVP/SBSP router where the transparent QoS control and forwarding mechanism is indicated. Figure 6 (right) shows the QoS control path for the integrated GS/SBS architecture over an r-hop path or VC. The user or DB component accesses GS/SBS at well-defined SAPs which, in the case of GS, goes through admission control and RSVP, and in the case of SBS, just goes through SBSP.

The principal desirable features of SBS include absence of resource reservation and admission control, well-defined simple user/network interface, per-flow QoS control with stateless (per-flow) routers, shield both user and network from complex computation responsibilities, distributed QoS control, small fixed-size headers independent of hop count, fine-granular usage pricing, zero-sum fair share, and user optimality via selfishness emulation and least cost service.

3.5 Proposed Research and Experiments

3.5.1 Implementation of Integrated GS/SBS Architecture

• Purdue-Wide GS/SBS Network The first and primary task will be to implement the integrated GS/SBS QoS provision architecture as a Purdue-wide internetwork. SBS, and a prototype version of GS, have been designed and tested using large scale simulations over WAN environments for VBR/CBR traffic with diverse QoS requirements. Figure 7 (left) shows a benchmark instance for vBNS (NSF's high-speed backbone network, the postcursor to NSFNET) carrying a diverse set of traffic flows. Figure 7 (right) shows the performance results of SBS with respect to the end-to-end QoS requirements of 11 application flows and the QoS delivered by the network. The guaranteed traffic flows (i.e., GS) are not shown as their QoS requirements are satisfied via reservation and admission control. Further details can be found in [5, 6, 27, 26].

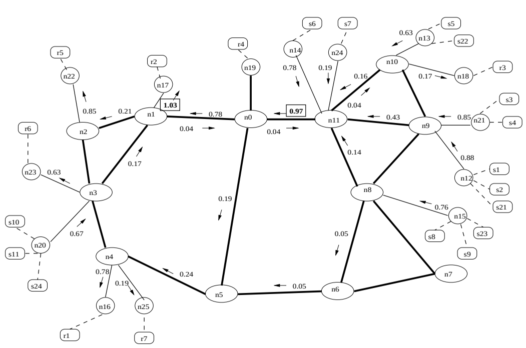

flow	pls. req.	delay req.	SBS		
			pls.	delay	cost
1	.05	.100	.0365	.031	49.64
2	.05	.100	.0341	.019	39.66
3	.05	.100	.0338	.041	59.66
4	.01	.080	.0002	.040	60.00
5	.01	.080	.0003	.045	60.00
6	.05	.100	.0341	.034	49.66
7	.01	.080	.0004	.008	30.00
8	.05	.100	.0363	.032	49.65
9	.01	.080	.0000	.043	50.00
10	.05	.100	.0264	.023	29.74
11	.01	.080	.0000	.045	60.00

Figure 7: Left: Prototype integrated GS/SBS architecture designed and tested using simulation on NSF's vBNS WAN topology. Right: Performance results for SBS with diverse application QoS requirements.

We will implement a corresponding architecture for a Purdue-wide QoS-sensitive campus network with the research coordination point anchored in the Computer Science Department. We will use the 6 GPS/RSVP-capable Cisco 72xx routers as the main backbone switching points with the 2 FORE ASX-1000 ATM switches acting as borderline aggregation points. The 8 FORE ASX200BX workgroup switches and the various GigabitEthernet, wireless LAN, and Myrinet switches will be used as intra-departmental transmission media for local transfer and processing of real-time and stored multimedia data.

• QoS Measurement Collection and Dissemination We will collect measurement data—both network traffic and user-specific QoS-related behavioral data—and use the DB component to store the data for later dissemination by our own research groups as well as the larger research community. This is similar in spirit to the NSF-funded measurement project INDEX being conducted at UC Berkeley.

However, whereas INDEX measures QoS data for end users dialing in from home using ISDN lines who use the same applications that a typical 28.8kb/s-limited user would—e-mail, telnet, ftp, and web browsing which are bandwidth-limited by both the ISDN lines and the current low-bandwidth world-wide Internet—our data will be significantly more comprehensive and reflect state-of-the-art multimedia applications including various forms of real-time and stored video for distance learning, remote sensing, and digital libraries to mention a few. As such, it will complement the measurement study provided by INDEX in an operational high-bandwidth internetwork environment.

• Large-Scale WAN Experiments It would be naive to think that the integrated QoS provision environment proposed in the project would give all the answers to the pressing problem of facilitating QoS in large-scale internetworked environments. Three of the PIs—Delp, Korb, and Park—are members of the Purdue-wide vBNS Advisory Board which oversees research projects conducted over Purdue's NSF-funded (NSF ANI-9729721) access to vBNS and Purdue's planned access to Abilene, a separate IP-over-SONET based Gigabit wide area network.

Several vBNS-based research projects involving QoS-sensitive transport of video and audio are in progress by Purdue researchers including multimedia information dissemination, QoS-sensitive transport of compressed real-time MPEG video/audio and self-similar traffic using adaptive control protocols (Park, NSF ANI-9714707), and network middleware experiments involving AT&T's Geoplex which is carried out across a number of vBNS-connected member institutions (Spafford, AT&T Labs).

Two additional members of the Computer Science Department with NSF-supported networking research projects who will join the integrated large-scale WAN experiments facilitated by GS/SBS, vBNS, and Abilene are Bhargava who will study adaptive QoS issues over high-bandwidth network environments which will extend current measurements carried out over the commodity Internet, and Szpankowski (NSF NCR-9415491, CCR-9804760) who will study the performance of pattern matching compression algorithms for video and audio when interfaced with a nonstationary network environment for variable rate transport.

3.5.2 Performance Comparison with Diff-Serv/Int-Serv Architecture

There are several proposals for differentiated service provision using aggregate flow control being discussed under the auspices of IETF's Diff-Serv Working Group. The GS part of the GS/SBS architecture is consistent with respect to its framework—resource reservation and admission control—with the schemes proposed in the Int-Serv Working Group of IETF which is based on RSVP.

Two of the Diff-Serv proposals—Clark's Assured Service and van Jacobson's Premium Service—approach the differentiated services problem from a different framework than that of SBS. The SBS (Park) approach has a number of desirable properties that the aforementioned schemes lack [5, 6, 26, 27] and it is interesting to provide a comparative study based on common benchmark experiments that delineate their performance differences. Both Clark and van Jacobson's schemes can be implemented using the same networking substrate as GS/SBS—RSVP-capable GPS routers (actually, Clark's scheme additionally requires a form of RED gateways which are supported by Cisco 72xx series routers)—and as such do not entail significant additional implementation effort since their central logic can be captured by end-to-end policing/marking controls.

Using the same context, we will study the issue of *deployability* in the larger vBNS/Abilene-based Internet2 context—i.e., not all routers may be GPS, RSVP, RED, or SBSP compliant—where the performance impact of noncompliant routers on end-to-end QoS provision is discerned and evaluated. This will also be tied to studies of user/transport-level QoS protocols over GS/SBS which may contribute toward masking the detrimental influence introduced by non-compliant routers using end-to-end control.

3.5.3 Refined End System Scheduling

An increasing number of applications use multithreading to enable asynchronous event-driven processing and the potential benefit of latency hiding that scheduling can bring. Whereas our CPU scheduling techniques enable threads to make guaranteed progress in an adaptive manner [34], new interprocess communication (IPC) mechanisms are needed to extend thread level guarantees to cross domain computation. To this end, research at the Software Systems and Architecture Lab (Yau, NSF ESS-9806741) has produced a *train* abstraction for predictable IPC performance [35]. Trains aim to decouple threads as a scheduling entity from their associated processes, and as such, enable threads to leave an existing process and enter a new one through a well defined *stop* exported by the latter. Experimental results show that trains allow server computation to proceed at the rate requested by their clients, and therefore are suitable for QoS.

3.5.4 User/Transport-Level QoS Control over GS/SBS

The GS/SBS QoS architecture cannot be expect to solve all QoS provision problems. Indeed, one of the primary motives of the proposed project is to investigate the dependencies between DB, networking, and security components as they relate to QoS and reflect user-specific customizations and needs into the networking part of the QoS provision architecture.

• WAN Experiments for Adaptable QoS We propose to conduct a series of experiments using various tools developed at the RAID Lab including WANCE and AVC that will aid in the development of policies for adaptability at the application, system, and network layer in the context of meeting quality of service requirements.

We will focus on three fundamental research problems: mechanisms for quality of service provision and control from the network layer based on the application requests; policies for secure, efficient and reliable transmission of multimedia data; and expert rules/guidelines for adaptability and its impact on various applications. Our research will focus on experiments and empirical measurement data collection involving advanced network services, protocols, techniques and their evaluation [3, 18].

• Real-Time End-to-End QoS Control Our project on self-similar traffic control (Park, NSF ANI-9714707) comprises of several subareas, two of which—congestion control for self-similar network traffic and real-time traffic control for bursty traffic using adaptive FEC—complement the SBS architecture by potentially making its services more robust via application/transport layer end-to-end QoS control.

We have shown that correlation structure present in long-range dependent traffic can be on-line detected and exploited for proactive feedback congestion control resulting in improved throughput. Our multiple time-scale predictive congestion control mechanism has been implemented and benchmarked for a number of congestion control protocols including TCP Reno and Warp Control [25], and in the case of TCP, performance improvements up to 60% have been demonstrated. We have also shown that user-specified QoS-sensitive transport of real-time traffic can be facilitated *end-to-end* over best-effort network environments (e.g., IP and ATM AAL5) using adaptive control protocols, in particular, adaptive forward error correction (AFEC). We have implemented AFEC for both real-time MPEG video, and more recently, MPEG audio transport, running under Solaris UNIX on Sun UltraSparc workstations. In the context GS/SBS, we will use our QoS-sensitive adaptive end-to-end protocols to provide enhanced robustness with respect to achieving a desired target QoS, especially, over SBS.

• Variable Compression Control Our work in pattern matching compression (Szpankowski, NSF NCR-9415491, CCR-9804760) uses tools from information theory and design of algorithms on words to provide effective algorithms for compressing video, images, and audio data [22]. The approach is based on a lossy form of Lempel-Ziv where using approximate string matching—in the time domain—we are able to achieve effective compression which is also grounded in results from information theory.

One extension that we seek to explore is varying the compression ratio of compressed VBR traffic as a function of network state such that a desired data rate is achieved. For example, when the network is congested, reducing the data rate is conducive to transmitting the video/audio information reliably and timely over the network. Conversely, when the network is underutilized or idle, then a high data rate may be tolerated by the network without causing significant packet drops or delay.

4 Security

The following sections summarize research activity related to building a QoS-sensitive security architecture encompassing several labs—COAST Lab (Spafford), CERIAS (Spafford), SERC Lab (Mathur), and Network

Systems Lab (Park)—at the Computer Science Department, where CERIAS (Center for Education and Research in Information Assurance and Security) and SERC (Software Engineering Research Center, NSF) are Purdue-wide and nation-wide centers, respectively.

4.1 Motivation

Security services, especially as they pertain to network security, are an integral component of any operational multimedia support infrastructure be they from the most basic aspect of access control to encryption and digital signatures. Mechanisms for achieving the principal security services—confidentiality, integrity, authentication, access control—are well understood, and in the network security context, can be implemented, if so desired, purely *end-to-end* without internal network support [32, 31].

Although the need for security and its import in any networked service provision architecture including MSI is evident and beyond dispute, there is significant apprehension among networking architects and OS designers to incorporate security services full-fledged into their respective architectures—even if part of the design, they are oftentimes not *implemented*—due to its immense overhead cost. Overhead and disturbances introduced both end-to-end and internally to the network can significantly impact the provisioning of QoS, and more generally, the reliability and trust of the overall system, and its influence must be explicitly incorporated into the network architecture to achieve effective QoS.

4.2 QoS-Sensitive Security Architecture

4.2.1 Security Mechanisms

The faithful rendering of security services encompasses two separate functional areas, one, in the delivery of security services related to the transport of *user data*—confidentiality, integrity, authentication, access control—and two, in the execution of analogous services for message flows that comprise *control* or *management information* such as connection set-up, routing, link-based flow control, and remote device management.

• User Plane Protection First and foremost, as far as the end user is concerned, a security architecture must provide confidentiality, integrity, authentication, and access control to message flows that stem from activities related to the transport of user data[1]. It is a well-known fact that all four services can be implemented end-to-end, i.e., as protocols running at end stations, without special support from the network medium itself.

Our security architecture, as it pertains to user plane services, will implement the four security services end-to-end in a *security middleware layer* (SML) acting between the transport and application layers of a typical 4-layer protocol stack (i.e., the ISO/OSI 7-Layer Reference Model minus the session and presentation layers). SML is implemented as a kernel-level library and, as such, is a "floating layer" that can be invoked from any layer in the protocol stack with the output passed to the designated layer—below or above. This is depicted in Figure 8 (left).

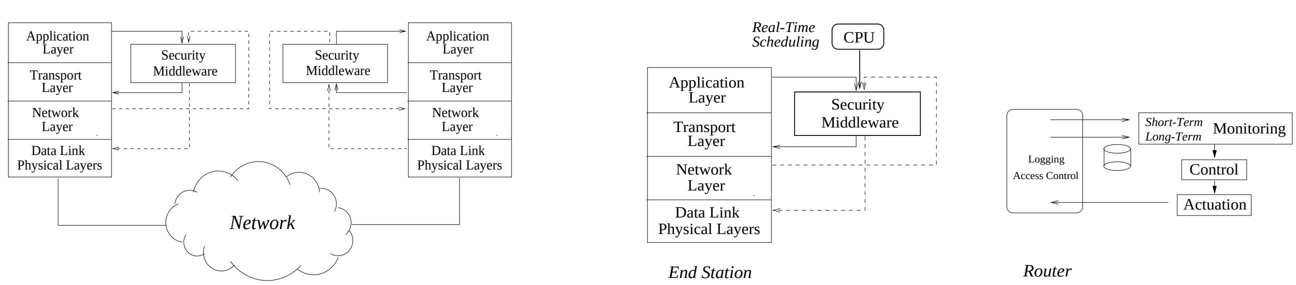

Figure 8: Left: 4-layer protocol stack augmented with security middleware layer implemented as a library. Right: Secure QoS mechanism at end station (real-time CPU scheduling) and router (adaptive security).

The security SAP is primarily accessed by the network component, and exports security services and their specific instantiations (e.g., encryption may be performed by a number of different algorithms). This selectivity facilitates flexibility and efficiency. When specialized support hardware (e.g., encryption chip) is available, device

[1]By no means is this list meant to be exhaustive; e.g., for anonymity, servers that implement a form of information hiding—many-to-one identity mapping—may be employed. However, the interaction between the user and servers is again end-to-end, and reduces to the original set-up.

drivers invoked from SML will interface with the hardware modules, transparent to the other layers in the protocol stack.

• Control/Management Plane Protection Control and management functions produce message flows which—modulo their *interpretation*—are just packets carrying bits and bytes, and thus can be protected end-to-end in the same manner as with message flows in the user plane. The only difference is that control and management protocols tend to reside at lower layers of the protocol stack—transport, network, and data link—but otherwise the protection requirements are exactly the same. The floating layer implementation of SML allows its services to be invoked from anywhere in the protocol stack (depicted in Figure 8 (left)) and thus facilitates efficient implementation.

Whereas management functions (e.g., SNMP for IP networks), for the most part, are end-to-end, control functions (e.g., PNNI for ATM networks) tend to be point-to-point, hence the cost associated with invoking security services increases proportionally with the hop count or path length over which a control message is propagated. Since activities such as setting up or tearing down of virtual circuits need to be performed frequently and expediently, imparting security to control functions induces a security-performance trade-off, and where in the trade-off curve to operate is dependent on several factors including organizational policies as well as individual user needs.

4.2.2 Secure QoS Mechanisms

• User Plane Services Given the selective security service invocation provided by SML, the two primary mechanisms by which QoS-sensitivity is achieved are *real-time scheduling* and *data rate prediction*. When security services are accessed through SML, its impact on QoS is similar to the impact that a MPEG video encoder (or decoder) would exert on the computational resources—i.e., CPU—of an end station. Namely, without adequate CPU allocation, even if the network were to deliver MPEG video frames in a timely manner, an end station may not be able to decode MPEG frames sufficiently speedily for real-time display. Thus depending on what security services are invoked and their specific instantiations, the scheduling guarantees requested of the real-time scheduler need to be changed accordingly.

An analogous correspondence holds for data rate prediction where, as in MPEG encoding, the ultimate data rate experienced by the network will depend on the security services selected due to the latter's data amplification effect. The bandwidth requested of the network, when establishing a QoS association, needs to reflect this data rate amplification factor so that the network can schedule its resources effectively.

• Control/Management Plane Services In general, the control plane and devices that partake in it such as routers are direct targets of attack due to their exposed nature (e.g., end stations can be shielded behind firewalls but gateways and routers in public switched networks are, by definition, accessible to rogue flows) and thus, if compromised, can have a significant impact on the functioning of the entire network including its exported QoS. Protection of the internals of a network from potentially debilitating attacks must be provided to facilitate stable, reliable services over the long run.

Proactively protecting all exposed network devices, although technically feasible, is not a viable solution due to its immense overhead and negative performance impact given that routers are expected to switch at Gbps speeds. The approach we adopt is called *adaptive security* where we minimize the on-line protective capabilities of a router—principally confined to access control and logging of audit trails—and reactively (at smaller time scales), and retroactively (at larger time scales), affect actions that curtail the impact of various attacks including intrusion, "correct" the distrupted state bringing it to a valid state, and take punitive actions that impart a significant cost to the attacker so as to achieve a deterrent effect.

4.3 Proposed Research and Experiments

The primary task will involve the implementation of the secure QoS architecture which is comprised of two phases, one, implementing QoS-sensitive user-plane security services using real-time scheduling, and two, realizing a QoS-sensitive—more generally, reliable and trusted—control/management plane which has a direct impact on performance.

4.3.1 Processing Latency Hiding

In the first phase, we will implement QoS-sensitive user plane services as a security middleware layer (SML) that is accessible by SAPs from the application layer—but also due to its library-based implementation—from lower layers in the protocol stack. QoS-sensitivity is imparted using real-time scheduling at end stations which will—when feasible—offset or hide the overhead introduced by processing needs of security services. The main problem to be solved in this context is the same problem that arises in scheduling of multimedia applications with real-time constraints; in the case of MPEG video, type (MPEG-I, II, IV, VII) and desired frame rate exert different computational loads on a host system's CPU and must be accommodated using a form of conversion which is part of QoS translation. We will omit the real-time scheduling description due to space constraints and refer the reader to the Networking section.

4.3.2 Adaptive Security

In the absence of attacks or other security related disturbances, the system achieved in the first phase will function properly—i.e., faithfully render user-plane security services—without adversely impacting scheduled or exported QoS. In the presence of intrusions, denial-of-service attacks, and other forms of malicious disturbances, the aforementioned system—although still functioning properly with respect to its security services—may be subject to disruptions that may cause it to fail in preserving its scheduled QoS. In networked systems of nontrivial size where *physical security* may be impractical to exercise, security related attacks loom as a potentially disruptive factor that needs to be actively managed.

The three principal components of adaptive security are monitoring, control, and actuation. The first monitors the system—at both small and large time scales—so as to detect various forms of anomalies and state perturbations. A prototype of such a distributed monitoring facility for intrusion detection called AAFID has been implemented at the COAST Lab (Spafford, NSA) in Perl.

The monitored information is collected and processed by secure QoS control algorithms at management stations—both on-line at small time scales and off-line at large time scales. The off-line component uses audit trail analysis to identify specific activities and the identity of intruders (Atallah & Spafford, CERIAS). Watermarking is used, in part, to check the authenticity of certain logged information and source tracing (Wagstaff & Atallah, NSA). The on-line component affects actions of a more responsive nature such as rerouting of traffic flows so that QoS is minimally perturbed by attacks and disturbances.

A combined monitoring and control system based on SNMP is being developed at the COAST Lab (Spafford & Park, Sprint). Reliability of composed systems—including its estimation and analysis—in the distributed systems context is being studied using hierarchical or renormalization techniques at the SERC Lab (Mathur & Park, SERC).

4.3.3 Secure QoS Experiments

At the end of the first phase of implementation, we will test the functioning of user-plane security services with respect to its performance impact by verifying that real-time scheduling with adaptive rate control and QoS translation is able to hide the processing costs and achieve complete transparency with respect to exported QoS.

During the second phase of the project, we will use the comprehensive reservoire of intrusion, penetration, and firewall testing tools available at the COAST Lab and CERIAS to perform various forms of attacks and test the architecture's resiliency at rendering QoS-sensitive security services under duress. We will test the effectiveness of using audit trails and watermarking to identify anomalies in the logged system activities and trace the identity of the intruder to its source. We will also perform sensitivity analysis at smaller time scales and measure the responsiveness of the adaptive security system to local disturbances. The Secure QoS Architecture will interact with the DB component when distributing its log files to storage devices and management stations dispersed over the network.

References

[1] S. Baqai, M. F. Khan, and A. Ghafoor. Multimedia communications. In W. I. Grosky, R. Jain, and R. Mehrotra, editors, *The Handbook of Multimedia Information Management*. Prentice-Hall, 1997.

[2] P. B. Berra, C. Y. R. Chen, A. Ghafoor, C. C. Lin, T. D. C. Little, and D. Shin. Architecture for distributed multimedia database systems. *Computer Communication (Special on Multimedia Communication)*, 13(4):217–231, April 1990.

[3] B. Bhargava, S. Li, S. Goel, C. Xie, and C. Xu. Performance studies for an adaptive video-conferencing system. In *Proceedings of the International Conference on Multimedia Information Systems (MULTIMEDIA 96)*, pages 106–116, 1996.

[4] J.-Y. Chen, C. Taskiran, E. J. Delp, and C.A. Bouman. Vibe: A new paradigm for video database browsing and search. In *Proceedings of the 1998 IEEE Workshop on Content-Based Access of Image and Video Database*, June 21, 1998.

[5] S. Chen and K. Park. A distributed protocol for multi-class QoS provision in noncooperative many-switch systems. In *Proc. IEEE International Conference on Network Protocols*, pages 98–107, 1998.

[6] S. Chen and K. Park. An architecture for noncooperative QoS provision in many-switch systems. To appear in *Proc. IEEE INFOCOM '99*, 1999.

[7] Y-T. Chen, R. L. Kashyap, and A. Ghafoor. Data placement for large read-only interactive multimedia information systems on multi-disk environment. In *Proc. of the SPIE/IS&T Symposium on Electronic Imaging Science and Technology*, Conference on Storage and Retrieval for Image and Video Database, January 1993.

[8] R. L. Cruz. Quality of service guarantees in virtual circuit switched networks. *IEEE J. Select. Areas Commun.*, 13(6):1048–1056, 1995.

[9] A.K. Elmagarmid, M. Rusinkiewicz, and A. Sheth, editors. *Management of Heterogeneous and Autonomous Database Systems.* Morgan Kauffmann Press, 1998.

[10] C. M. Fiduccia and R. M. Mattheyses. A Liner-Time Heuristic for Improving Network Partitions. In *Proceedings of the 19th ACM/IEEE Design Automation Conference DAC 82*, pages 175–181, 1982.

[11] D. Le. Gall. Mpeg: A video compression standard for multimedia applications. *Commun. of ACM*, 34(4):46–58, April 1991.

[12] H. Jiang and A. K. Elmagarmid. Spatial and temporal content-based queries in hypervideo databases. *Special Issue on Multimedia Data Management, The Very Large Database Journal*, 1998.

[13] H. Jiang and A. K. Elmagarmid. Wvtdb - a web-based videotext database system. *IEEE Transactions on Data and Knowledge Engineering*, November 1998.

[14] H. Jiang, A. Helal, A. K. Elmagarmid, and A. Joshi. Scene change detection techniques for video database systems. *Journal of ACM Multimedia Systems*, 1998.

[15] T. Johnson and S. Prabhakar. Tape group parity protection. In *Proceedings of the 16th IEEE Symposium on Mass Storage Systems*, San Diego, CA, March 1999.

[16] M. F. Khan and A. Ghafoor. Design and evaluation of disk-scheduling policies for high-demand multimedia servers. *To appear in the Proceedings of IEEE International Conference on Data Engineering, Sydney, Australia*, March 1999.

[17] W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2:1–15, 1994.

[18] Shunge Li and Bharat Bhargava. Active Gateway: A Facility for Video Conferencing Traffic Control. In *Proceedings of COMPSAC'97, Washington, D.C.*, pages 308–311. IEEE, August 1997.

[19] T. D. C. Little and A. Ghafoor. Network consideration for distributed multimedia object composition and communication. *IEEE Network (Special Issue on Distributed Applications for Communications)*, 4(6):32–49, November 1990.

[20] T. D. C. Little and A. Ghafoor. Spatio-temporal composition of distributed multimedia objects for value-added networks. *IEEE Computer (Special issue on Multimedia Information System)*, 24(10):42–50, October 1991.

[21] T.D.C. Little and A. Ghafoor. Multimedia synchronization protocols for broadband integrated services. *IEEE Journal on Selected Areas in Communications*, 9(9):1368–1382, December 1991.

[22] G. Louchard and W. Szpankowski. On the average redundancy rate of the lempel-ziv code. *IEEE Trans. Information Theory*, 43(2–8), 1997.

[23] B. Özden, R. Rastogi, and A. Silberschatz. A disk-based storage architecture for movie on demand servers. *Information Systems, Special Issue on Multimedia Information*, 20(6):465–482, 1995.

[24] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the multiple node case. *IEEE/ACM Trans. Networking*, 2(2):137–150, 1994.

[25] K. Park. Warp control: a dynamically stable congestion protocol and its analysis. In *Proc. ACM SIGCOMM '93*, pages 137–147, 1993.

[26] K. Park. SBS: A global scalable network architecture for stratified quality of service provision. Technical Report CSD-TR-98-040, Department of Computer Sciences, Purdue University, 1998.

[27] K. Park, M. Sitharam, and S. Chen. Quality of service provision in noncooperative networks: heterogeneous preferences, multi-dimensional QoS vectors, and burstiness. In *Proc. 1st International Conference on Information and Computation Economies*, pages 111–127, 1998.

[28] K. Park and W. Willinger, editors. *Self-Similar Network Traffic and Performance Evaluation.* To be published by Wiley Interscience, 1999.

[29] S. Prabhakar, D. Agrawal, A. El Abbadi, and A. Singh. Scheduling tertiary I/O in database applications. In *Proc. of the 8th International Workshop on Database and Expert S ystems Applications*, pages 722–727, Toulouse, France, September 1997.

[30] K. Shen and E. J. Delp. Wavelet based rate scalable video compression. *IEEE Transactions on Circuits and Systems for Video Technology*, 11, November 1998.

[31] E. Spafford and S. Kumar. A software architecture to support misuse intrusion detection. In *Proceedings of the 18th National Information Security Conference*, October 1995.

[32] E. Spafford and C. Schuba. A reference model for firewall technology. In *Proceedings of teh 13th IEEE Computer Security Applications Conference*, pages 133–145, December 1997.

[33] J. Stankovic, S. H. Son, and L. Jorg. *Real-Time Database and Information Systems: Research Advances*, chapter BeeHive: Global Multimedia Database Support for Dependable, Real-Time Applications, pages 409–422. Kluwer Academic Publishers, Norwell, MA, 1997.

[34] D. Yau and S. Lam. Adaptive rate-controlled scheduling for multimedia applications. *IEEE/ACM Transactions on Networking*, 5(4), 1997.

[35] D. Yau and S. Lam. Migrating sockets – end system support for networking with quality of service guarantees. *IEEE/ACM Transactions on Networking*, 6(6), 1998.

RESULTS FROM PRIOR INFRASTRUCTURE GRANTS

The previous infrastructure grant for our department was SoftLab, CDA 9123502, which expired in September after a 1-year no-cost extension. The PIs were Christoph Hoffmann, Elias Houstis, and John Rice. The objective of the grant was to (a) develop leading-edge experimental research programs establishing software laboratories, and (b) establish a graduate program in computational science. The accomplishments for both are discussed next. In the grant period, Hoffmann published 1 book, 8 book chapters, 12 refereed journal papers, 5 reviewed conference papers, and 8 invited papers. Houstis and Rice also published 2 books, 4 book chapters, 29 journal papers, 20 refereed conference papers and 14 other conference papers which are directly related to the SoftLab project.

A. EREP WORK (Hoffmann). A major research component has been the investigation of computer-aided design for (discrete) manufacturing, under the direction of Christoph Hoffmann. The work has developed a high-level, editable representation (Erep) that supports the recent design paradigms of feature-based, constraint-based design. The project implemented a full CAD system as test bed for the ideas, including a variational constraint solver, a Motif-based graphical user interface, a sketcher, and a design compiler in charge of geometry creation from the high-level description of the design process. Much of the work has been supported by the Office of Naval Research and by other NSF grants. The system runs under X/Solaris and under Windows/NT.

The impact of this work has been significant. The representational ideas were leveraged in the industry-led ARPA ENGEN project, managed by the South Carolina Research Authority. It was recognized that Hoffmann's Erep provided a neutral description for features and constraints, and thus was a good candidate for devising extensions to the international STEP standard for product data exchange. Participants in the project included Ford, SDRC, Computer vision, and Parametric Technology.

The Erep work also led to the Open CAD Architecture for Interoperability (OCAI) initiative undertaken by the National Center for Manufacturing Sciences (NCMS). This project, with participation over several years by NCMS member companies, developed a master model architecture. Again, the work recognized the significance of the technical aspects of the Erep, in particular, the ability to persistently associate attributes and relationships with elements of changing product shapes.

Geometric constraint solving is an important aspect of the Erep work. There has been substantial progress in this component research as well, including how to deal with conics and cubic curves, with spatial constraints, how to integrate symbolic constraints, and a highly efficient constraint decomposition algorithm based on network flow.

B. PSE WORK (Houstis and Rice). The project defines a problem solving environment (PSE) to solve partial differential equation (PDE) problems. The structure is to encompass the entire solution process ranging from very low-level modules to complete problem solving environments, using an object-oriented approach with multiple levels. At the top level are large, complex objects like PDE PROBLEM, NUMERICAL PDE PROBLEM, SPECIFICATIONS, SOLUTION. At the bottom level are utility objects like the BLAS (Basic Linear Algebra Subroutines) to manipulate vectors and arrays. Each object has substantial identifying information about its type, capabilities, properties, and interface requirements. This allows the structure to be open ended. Within this framework a set of problem solving kernels is defined (e.g., for discretizing PDEs, meshing domains, solving linear systems, plotting results) and a prototype, scalable subset is being constructed. This approach allows the community of PDE experts to exchange software, to build huge systems more cheaply, and to reuse software components in parallel applications. A world wide web server for the //ELLPACK system has been set up, where the user interface operates on a local machine and the PDE solution runs on other machines on the network. See `http://pellpack.cs.purdue.edu`.

We have developed a new paradigm and systems for the simulation of complex and composite physical objects based on PDE models. A global/local software framework is provided capable of capturing these interactions and iterating until a global solution is obtained. This paradigm is embedded in a very high level graphical user interface implemented with visual programming techniques, object-oriented software modules and computational agents. A Repository of "Legacy" PDE Software has been set up. It contains all the state-of-the-art public domain or "almost" public domain PDE solving software; see
`http://www.cs.purdue.edu/research/cse/pdelab/pdepack/pdepack.html`.

C. INTELLIGENT COMPUTING (Houstis and Rice). A Scientific problem solving environment requires substantial "intelligence" if it is to be used by ordinary scientists and engineers. This aspect of SoftLab led to two designs for a software recommender system and one of them, *Pythia* has been implemented. This work was focused on second order elliptic partial differential equations and incorporated much of these PIs previous work on the performance analysis for these equations. In addition, a number of studies were made of the decision making methodologies to be used, e.g., case-based, neural net, neuro-fuzzy, and inductive logic programming. An application of Pythia was made to the area of numerical quadrature to demonstrate that the methodology and systems is not constrained to PDE problems.

D. APPLICATIONS (Houstis and Rice). In collaboration with colleagues in Chemical Engineering, two problem solving environments were designed and created: (1) *BioSeparationLab* for the control of processes that purify foods and drugs, and (2) *MicroElectronicsLab* for modeling the process of CVP (chemical vapor deposition) on the surface of microprocessor chips. BioSeparationLab is in production use now.

E. THE PURDUE COMPUTATIONAL SCIENCE & ENGINEERING PROGRAM. An interdisciplinary graduate program in Computational Sciences and Engineering has been established, supported by three new computer science courses:

501 – Introduction to CS&E
525 – Parallel Computing
530 – Scientific Visualization

The program has been set up as an interdisciplinary graduate program with broad participation by other departments on campus. We have hired a two new CS&E faculty , Elisha Sacks and Ananth Grama, as part of Purdue's cost sharing for SoftLab.