

ISSUES AND TECHNIQUES IN NETWORK-BASED DISTRIBUTED HEALTHCARE: QUALITY OF SERVICE AND BANDWIDTH

Ying-Lung Liang & Stan McClellan*
Center for Telecommunications Education & Research
Electrical & Computer Engineering
Univ. Alabama at Birmingham
Birmingham, Alabama, 35294
Email: smcclell@uab.edu

Ken Burst
Broadband Laboratory
BellSouth Technology Assessment Center
Bellsouth Telecommunications
Hoover, Alabama 35244
Email: kburst@btac.com

ABSTRACT

With the advent of the “next generation Internet” and various related technologies, we observe important trade-offs between emerging network capabilities and design related application requirements for network-based distributed healthcare systems. One of the aspects of these emerging network capabilities is *Quality of Service*, or QoS. QoS guarantees are a necessary characteristic of “next generation” networks which will have a profound impact on the deployment of advanced, network-sensitive medical applications. In this paper, we discuss some of the underlying aspects of QoS technologies and possible considerations in designing network-based medical systems.

INTRODUCTION

The essence of the Internet2 consortium (<http://www.internet2.edu>) is to accelerate the deployment of “next generation” connectivity technologies. The fundamental concept that distinguishes Internet2 from the best-effort nature of the commodity Internet is the proposal of some capability for guaranteed Quality of Service (QoS). Various network technologies and topologies exist which are more-or-less “QoS-capable”, such as Asynchronous

Transfer Mode (ATM), or ethernet with virtual-LAN and class-of-service extensions (IEEE 802.1 p/q). However, there are several deterrents to near-term, generally-available deployments of QoS capabilities. Some of the most significant deterrents relate to the heterogeneous, multi-domain, production nature of the distributed Internet, and the lack of a robust approach to ensuring QoS requests across domain boundaries. Additionally, the lack of a unified and widely accepted application-level interface for specifying QoS requirements prevents active participation by applications and users in the complete process. These deficiencies are being addressed by the Internet2 “Quality of Service Working Group” and the IETF “Differentiated Services Working Group”, among others. Long-term solutions which are both robust and implementable may be forthcoming based on these efforts.

The notion of “Quality of Service”, much like beauty, is “in the eye of the beholder.” As such, it is almost impossible to separate the deployment of QoS capabilities from a thorough evaluation in the context of an application (or class of applications). Unfortunately, since application-level requirements are many and highly variable, a thorough evaluation of QoS technologies in the context of *all* applications is impossible. Here, we propose compromises: (1) evaluation of potential QoS mechanisms jointly in the context of physical network elements and a representative class of “next generation” applications, and (2) the development of mechanisms which end-users can access to participate directly

*Corresponding Author. This paper was prepared in affiliation with David A. Conner, David Green, Gregg Vaughn, Jay Goldman, Gary Grimes, and Murat Tanik of the UAB Center for Telecommunications Education & Research. The authors also wish to acknowledge the contribution of Harold E. Crane of ATI and Peter Adams of CTER.

in the negotiation between network capabilities and application requirements. Additionally, we propose the evaluation of these approaches in the context of “access technologies” whereby users & applications in the broader implementations of LAN & WAN technology (ie. not just LAN-connected users) can access QoS-capable network services. A key component to these approaches is the evaluation, via embedded simulation, of a “multiprotocol” (i.e. multiservice, heterogeneous) network infrastructure where the “common ground” is the suite of Internet protocols (IP).

“ACCESS” vs. “CORE” TECHNOLOGIES

The deployment of functional QoS capabilities is heavily dependent on “core technologies” such as interfaces between independent subdomains that act as peers to exchange routing and traffic management information. The evolution of QoS-capable core technologies has been underway for some time. Approaches such as RSVP, native ATM QoS, and others have been deployed, examined, and found to be individually lacking in several areas. The near-term “next steps” for this class of services is, in some sense, fairly well defined by the scope and magnitude of the task (McClellan *et al.*, 1999).

In addition to the dependence on “core technologies”, the thorough evaluation of “access technologies” is crucial in developing and validating a functional, broadly deployable QoS framework. In our paradigm, the terminology “access technologies” includes the interaction between application-level requirements, the network protocols that are present at the boundary between application and network, the logical topology of the local network, and the unique physical layer capabilities which are provided by various network architectures.

In developing a strategy for evaluating QoS mechanisms from the perspective of “access technologies,” we observe a few relatively obvious, but important typical situations. For example, (1) QoS requirements are expressed at the edge of the network, or at the point of ingress, prior to (or instead of) propagation through/between sub-domains, and (2) there are generally two distinct physical contexts through which ingress is accomplished. Here, we briefly discuss approaches to network architecture, evaluation of services, and general concepts which are being formulated to address the deficiencies of the commodity Internet for carrying “premium” traffic.

LAN-BASED

In the context of LAN-based ingress, users are primarily ethernet-connected hosts which may lie on an ATM core, use various isolation mechanisms (switching), have gigabit ethernet uplinks, and so on. Regardless of these conditions, the logical/physical topology of the network directly affects

the performance of the application. The types of QoS mechanisms deployed and the level of support these mechanisms enjoy at various “strata” in the network are significant concerns. These considerations are particularly important in the context of adequate, network-intensive healthcare applications. For example, the existing logical/physical infrastructures of an institution such as a large hospital or clinic may be structured in a manner which renders QoS-sensitive traffic impossible without significant upgrades.

In LAN-based networks, and in the concept of the “Next Generation Internet”, there tend to be relatively few general approaches to ensuring the integrity of data streams. These approaches tend to (a) shape the aggregate traffic stream with simplistic regard for the needs of individual streams, or (b) have the capability to ensure specific service characteristics for individual streams at the expense of tremendous complexity for the user and/or application. The “Class of Service” (CoS) approaches define a small set of generic behaviors, which if properly defined and applied, can approximate the actual service requirements of the data stream (McClellan *et al.*, 1999). Approaches capable of “true QoS” infer an amount of homogeneity in the underlying structure of the network, and exchange a large amount of complexity and effective throughput for highly specific constraints on multiple statistical parameters of the data stream.

The current structure of the Internet relies heavily on conventional forwarding devices (routers) which must store, examine, and transmit individual clusters of information (packets) from multiple simultaneous data streams. These actions are logically similar to the performance of a shipping clerk or production-line worker. Each item from the assembly line is inspected, wrapped, boxed, and loaded for transport. Broken items must be tossed aside, maximum efficiency is achieved when all items are “similar”, and the introduction of “special” items creates significant difficulties in resource allocation, etc. The introduction of data streams into the commodity Internet which have particular needs for QoS is logically equivalent to interlacing “special handling” items into a production line or assembly line. Depending on the nature of the data (size, shape, characteristics of the item), the ramifications of transport are felt along the entire production line. Tactical reconfiguration (or the capability for it) is absolutely necessary to maintain adequate delivery volume.

Shortcut Routing & Switching

The integration of “routing” and “switching” in the Internet seeks to balance the competing goals of packetized data transport and stream data transport in much the same way that a conventional assembly line can be adapted to support certain classes of “special” items. Switching technology for networking is closely analogous to a highly special-

ized assembly line, optimized with robotics and tight control structures to maximize throughput and minimize implementation cost. Switching architectures tend to be implemented in special-purpose hardware which can efficiently process highly structured data streams.

Clearly it is impossible (impractical) to implement individual, highly optimized production/assembly lines for each product in a repertoire of potential products. Because of this fundamental limitation, the fusing of “routing” and “switching” technologies has become a very important step in the evolution of multiservice network capability. The “switch-router” or “layer 3 switch” is an evolutionary step towards a two-stage production line wherein “forward” tactical information gathered from the data stream or topology is used to rapidly and temporarily reconfigure subsequent product manipulations. In short, shortcut routing in switch-router networks relies on the capability of a tactical entity from a particular domain to “cross over” into the sphere of influence of another domain. In this setting, a route is established and resources may be allocated across a wide-ranging switching infrastructure.

In the context of data networking, “flow detection”, “shortcut routing”, and “IP switching” are terms which are commonly used to describe this process. Typically, the efficient, synchronous data handling capability of a switch architecture is slaved to the slower, high-level processes which accompany traditional routing technologies and effect the distribution of network topology. The analogy of the resources available as “front lines” or “trenches” (ie. switching) being directed by “headquarters” (ie. routing) based on fed-back tactical information is quite good in this case, particularly when combined with the notion of “sealed orders” (or, shortcut paths) (McClellan *et al.*, 1999).

Differentiated Services

The current “large scale” approach to supplying some form of QoS in Internet2 lies in the approach recommended by the DiffServ Working Group (Sikora & Teitelbaum, 1999). This approach, which provides a mechanism for offering something *other than* best-effort service, is likely to function well for certain classes of traffic. However, additional research is needed in areas pertaining to application-initiated QoS requests, user-interface specifications, and “situational awareness” for users of the network. By aggregating “similar” flows into equivalence classes, the DiffServ approach hopes to create “virtual leased lines” without the need to maintain per-flow state or signalling information in all intervening packet-switching nodes (routers). By nature, this approach is primarily directed at “core” and interdomain issues. As such, most of the inherent complexity is pushed to the edge to exploit distinctions between “core” and “edge” data handling. In this fashion, examination of

per-flow characteristics is done *once* at ingress, with simple per-hop behaviors and aggregate policing performed by task-optimized core nodes.

To enable distributed system configuration (policing, leaf routers, etc.), each independently administrated network implements a “bandwidth broker.” The bandwidth broker (as yet not completely defined) acts as a proxy on behalf of a user or application to request traffic flow through transit networks. As a result, separately administered networks may be implemented with differing technologies, policies, and network management strategies.

The Internet2 QBone is a subset of Internet2-connected institutions who have volunteered to deploy and evaluate recommended QoS technologies. The initial DiffServ capability proposed for deployment in the QBone is the “Premium Service”, and is supposed to support flows up to a maximum rate with low loss and minimal queueing delay. Later deployments in the QBone will include “predictive” services with some ability to dynamically adapt, as well as “precedence” services with some guarantees for relative performance. Since this is an evolutionary deployment, there are several outstanding issues associated with Premium Service, such as:

- **Classification:** comparing a packet’s characteristics against existing agreements for service level,
- **Marking:** setting flags (Type-Of-Service bits) in the packet’s header to indicate subsequent eligibility for Premium Service,
- **Shaping:** a shallow token-bucket
- **Prioritization:** 2 queues in all forwarders with a goal to keep the high-priority queue empty,
- **Policing:** dropping packets when resources are unavailable.

At ingress, the Classification, Marking, Shaping, and Policing functions are performed by a “leaf router.” Conforming packets are placed into a high-priority queue for express-forwarding. This configuration is tremendously sensitive to misconfiguration at ingress since implied internal “trust domains” mean that internal forwarders *don’t* perform shaping, etc. Instead, internal “clouds” contract for simple, bilateral QoS profiles on aggregate traffic. Inside each cloud, local construction (ie. relative ethics) are applicable.

In the DiffServ model, a combination of assumptions regarding architectural and traffic characteristics are required. These include (1) strict priority forwarding at internal nodes, (2) no bursts allowed at the peak rate, and (3) heavy overprovisioning of network capacity are required to minimize latency and jitter (Nichols, 1999). As a result, open research topics include techniques for classification of packets, hardware & algorithms to ensure expedited forwarding at congestion points, and admission control & policing techniques.

User-Initiated QoS

Although promising and highly scalable, several aspects of DiffServ-based QoS mechanisms are unresolved, and performance for specific applications is questionable. Another potentially useful approach involves the concept of user-initiated QoS via “situational awareness.”

Although the terminology “situational awareness” derives from military requirements for highly integrated man-machine interfaces, it is useful to employ it in the context of advanced networking and application requirements. For example, a pilot in a military aircraft achieves “situational awareness” by assimilating multiple, simultaneous inputs from radar and other sensors. To be useful, this data must be parsed, analyzed, and prioritized before being presented to the user. Automated data reduction is necessary because decisions regarding weapons budgets, tactical maneuvers, etc. must be based on an instantaneous combination of summary data and personal experience. Similar considerations are involved in network-based medical applications.

The terminology “situational awareness” is also particularly applicable in other contemporary environments where human capacity for information gathering and interpretation is being (has been) exceeded. It is equally applicable in situations where personal interpretation of data is important in producing context-sensitive actions. For example, individual automobile drivers in large metropolitan areas derive situational awareness from TV/radio traffic reports. These reports may be processed differently by each driver depending on his/her instantaneous priorities, and corresponding actions may be initiated. The Georgia Dept. of Transportation has applied “next generation” network-based data gathering to this societal phenomenon by deploying cameras and other sensors along heavy traffic routes (<http://www.georgia-navigator.com>). As data is collected and minimally interpreted, a convenient web-based interface allows drivers to collect personally useful information about the status of roadways. This same concept is also present on many Internet software repositories. Often, users are given some statistics, location, or reliability information about software download sites in order to make informed decisions about the download process. In each of these scenarios, general information about the “state” of a public resource (roadways, traffic, network connectivity) is fused instantaneously with specific information about the “state” of a personally-important process (commuting, gathering software). Regardless of the scenario, the joint optimization between “general state” and “specific state” is a fundamental aspect of the performance of each task or process. This joint optimization is embodied in the Internet2 mantra: “it’s the application, stupid”; as well as the in the concept of application-initiated QoS.

The DiffServ proposition is essentially a non-optimal ap-

proach to solving this problem in the context of Internet networking. DiffServ nominates routers and “bandwidth brokers” as one-sided proxies to act on behalf of individual users and applications. These actions are taken with respect to a particularly limited set of average characteristics. Wholesale deployment of ATM technology is also essentially a non-optimal approach to the problem because it sloughs the entire complexity of optimization directly onto the user. To fully utilize the tremendous potential of an ATM infrastructure, individual users must perform joint data reduction between the specific information about a process (or collection of processes) and the state of the network. ATM signalling can alleviate some of the instantaneous burden of this process, but applications must still be characterized and modified to allow proxy-based negotiation with the network. Another flaw in reliance solely on ATM-based QoS guarantees is the “by design” heterogeneous structure of the distributed Internet. This reality violates the principal assumption of QoS via ATM mechanisms (McClellan *et al.*, 1999).

The underlying difficulty with these examples of reality-constrained deployment of technology is the use of a “layered” or “independent” approach to system integration/optimization. This layered approach has been used with some success in deploying classical IP routing in IP-over-ATM network topologies, but it remains an essentially piecewise optimization with many flaws. Unfortunately, while successfully using the strengths of various independent approaches in cobbling together a whole somewhat greater than the sum of its parts, much potential has gone unrealized.

A similar situation is likely with broad-based video distribution in next generation IP networks. Contemporary video encoding and transport mechanisms have been standardized and are widely understood (ie. H.263, MPEG, RTP, RTCP, etc.). Additionally, frameworks for point-to-point, point-to-multipoint, and gateway-based (transcoding) video distribution are available (ie. H.320, H.323, IP multicast, etc.). However, the “layered” use of these independently optimized technologies is substandard and can produce significant problems in network-based delivery as well as reception of video streams.

The best long-term solution to such problems is joint optimization between encoding, transport, and infrastructure components. In this case, the technological landscape shifts too rapidly for such an approach to be workable. A potential approach to approximate joint optimization of these distributed components is the deployment of user-level tools which provide “situational awareness” and “tactical feedback” to users who have immediate, vested interest in network utilization. Basic developments of this nature have been widely deployed on the Internet for control of “web cams” and the like. However, these technologies are *pas-*

sive in that they don't close the distributed feedback loop that simultaneously involves network topology, user request, and process requirements. Additionally, vendor-specific technologies exist for integrated observation and control of carrier-class telecommunication networks. These technologies present intuitive, user-level topology and connectivity information to network administrators, but typically don't include highly functional portals whereby individual users can affect the instantaneous virtual topology or end-to-end throughput characteristics of the network.

In this light, the systematic, application-driven development of user-level tools and middleware that transparently affect the virtual topology of the network is a topic of some interest.

WAN-BASED

In the context of WAN-based ingress, the logical/physical topology of the network is more likely to be based on circuit-switched paradigms, and is also likely to be much less flexible than LAN-based ingress. Additionally, the commercial interests of the Network Service Provider (NSP) combined with the availability of QoS-capable products as a competitive advantage may be very attractive from the standpoint of overall system optimization. As a result, NSP's may be willing to explore the optimization of potential services prior to deployment.

UAB's CTER has established a partnership with BellSouth Telecommunications, Inc. for the purpose of "in situ" evaluation of "next generation" network access paradigms. BellSouth's Technology Assessment Center (BTAC) in Hoover, AL is a unique facility where prospective services for the public network are rigorously evaluated prior to full release. To accomplish this level of "real world" exercise in a reasonable timeframe, all testing and evaluation projects are integrated with duplicated public network facilities, and vendor representatives are able to "lease" evaluation time on the BellSouth facilities prior to introduction of software/hardware changes, upgrades, installations, and so on.

The availability of numerous access technologies which can be used for testing/evaluation of "next generation" QoS, CoS, etc. proposals can provide significant insights into the applicability of various approaches for eventual commercial deployment. For example, several evaluation scenarios and approaches are briefly outlined below:

- Evaluation of technologies to dynamically allocate network services based on application-level events
- Focused testing to determine how potentially deployable public network architectures can meet the performance levels proposed for Internet2
- Optimization of topological deployments in the context

of various QoS parameters

- Focused testing to determine how advanced applications can take maximum advantage of commodity network architectures deployed for Internet2
- Development of applications which are structured to take maximum advantage of technologies such as xDSL with QoS
- Development and evaluation of algorithms and procedures for monitoring/billing network services based on application-level characteristics
- Development and evaluation of interfaces and/or "wrappers" or "containers" into which well-defined legacy applications can be placed in order to request/negotiate application-level QoS in the public network

Although these concepts may have merit, they all implicitly involve some knowledge of application-level requirements.

Results of WAN-class network testing

Regardless of network architecture or topology and irrespective of QoS mechanisms, a precise determination of the adequacy of network services for individual applications is very important. This is particularly true in the context of network-based medical applications.

Figure 2 shows representative application-specific results of QoS testing for a telemedicine system designed for remote pathology (McClellan *et al.*, 1996; Winokur *et al.*, 1996; Grimes *et al.*, 1997). This testing was conducted using a particular WAN-class ATM network, as shown in Figure 1. Notice from Figure 1 that 7 streams of various classes of ambient traffic are being generated by the ATM test set (Adtech AX/4000). These traffic streams pass through and completely fill an OC3 link (155 Mbps) which feeds into an optical splitter. The splitter duplicates the OC3 traffic six times, producing a total of 42 streams of ambient traffic at an aggregate rate of 930 Mbps. This entire traffic volume is mapped directly through the OC3 "trunk link" between the central switches. Additionally, the ATM LANs on either side of the WAN cloud are connected at OC3 level, and the traffic passing between LANs is directed through the "trunk link" using an ATM virtual path. Each traffic stream entering the trunk link is assigned a particular "traffic descriptor" which creates a hierarchy of traffic characteristics. Clearly, with the telemedicine traveling through the trunk link which is **heavily** oversubscribed, the quality of service characteristics for the application data streams is particularly "at risk".

The application data in this case was recorded on an isolated test network during clinical trials of the telepathology system, and so is representative of actual diagnostic sessions for such a system. As a result, this data and test scenario forms a useful "baseline" for evaluating QoS parameters such as latency, required bandwidth, and so on. In the clinical trials of the system, specimens were numbered so that

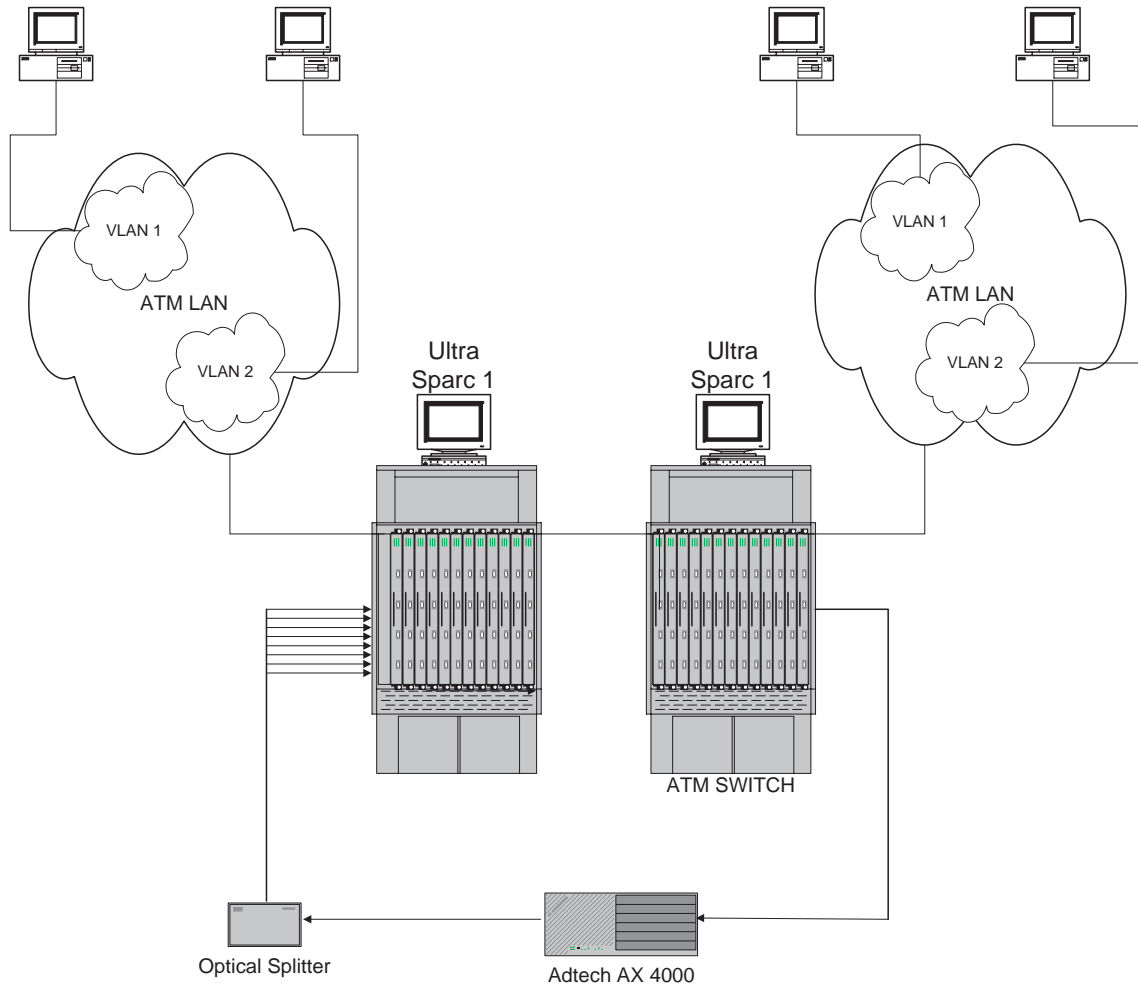


Figure 1. Test setup.

the study would be blind. As a result, the participating doctor's name and the slide number are shown in the plot. In the figure, the horizontal axis is cross-traffic intensity and is related to a variable-bit-rate stream with increasingly higher average rate and variability. The vertical axis is total session time in seconds. Each session was recorded exactly and iterated many times under the same network conditions, and a representative session is shown. As a result of operating system differences, network, and processing randomnesses, a session doesn't run exactly the same way twice. Events inside the session are slightly skewed, the complete session duration is different, and so on. In this application, the overall session duration must be kept to a minimum to properly enable diagnostic requirements.

At each intensity value, the circle with crosshairs is the average (mean) session duration for a particular traffic intensity (noise level). The vertical bars indicate one stan-

dard deviation about the mean. Notice that in all cases, the skewness introduced by system randomness is minimized. The circles are roughly at the same vertical offset (average session duration), and the spread of the distribution is quite small. Additionally, the mean and variance the distribution for each noise level don't increase as the cross-traffic load increases. In fact, each scenario is very similar to the case where noise level is zero, or the network is completely isolated.

In most cases, the standard deviation for each traffic intensity level is around 1.5 sec for "total session time", which can range from 200 to 700 sec and describes a complete diagnostic encounter. This is the most critical time in the telepathology scenario because a "session" starts when the remote pathologist begins to look at the tissue, and it stops when he reaches a diagnosis and surgery can resume. See (McClellan *et al.*, 1996) for details. The fact that essen-




Figure 2. Siegla, Slide 07, Study1.

tially no variation is present in this overall time indicates that even when the simulated sessions (repeated *many* times) are passing through the intervening WAN-class ATM cloud under very heavy ambient load, the traffic management of this system architecture *absolutely guarantees* transit to within a second or so over several hundred seconds.

In this case, based on matched traffic descriptors in a “passive” environment (ie. no ATM signalling, high-order IP routing algorithms, etc.), absolute QoS is demonstrated for an individual, medically-related datastream with highly specific delivery requirements. These results are interesting in that the adequacy of a particular end-to-end system for transporting real-time, interactive medical sessions has been validated in the context of the actual application.

CONCLUSION

QoS guarantees are a necessary characteristic of “next generation” networks which will have a profound impact on the deployment of advanced, network-sensitive medical applications. As such, it is almost impossible to separate the deployment of QoS capabilities from a thorough evaluation in the context of an application (or class of applications). Here, we have discussed many topics relevant to QoS in distributed, network-based healthcare systems. These top-

ics include the Internet2 DiffServ recommendations and a distributed testbed for validating these approaches in actual implementations (the QBone).

We have also proposed two approaches for ensuring user-validated QoS for highly sensitive applications: (1) evaluation of application requirements and QoS mechanisms jointly in the context of physical network elements, and (2) the development of mechanisms which end-users can access to participate directly in the negotiation between network capabilities and application requirements. In the spirit of the QBone, we verify the operation of a representative telemedical application in an embedded simulation of a “multiprotocol” (i.e. multiservice, heterogeneous) network infrastructure where the “common ground” is the suite of Internet protocols (IP) and QoS requirements are absolute.

ACKNOWLEDGMENT

This work was sponsored, in part, by grants from the National Institutes of Standards and Technology Advanced Technology Program (70NONB581011) administered by SCRA, by BellSouth Telecommunications, and by an equipment grant from Network Equipment Technologies (N.E.T.). This material is also based upon work supported by the National Science Foundation under Grant No. 9729500.

REFERENCES

- Grimes, G., McClellan, S., Goldman, J., Conner, D., Vaughn, G., Kujawski, E., Winokur, T., Fleming, W., & McDonald, J. 1997 (February). Applications of virtual reality technology in pathology. *In: Proc. Medicine Meets Virtual Reality 5*.
- McClellan, S., Winokur, T., & Kujawski, E. 1996 (October). A prototype system for telepathology. *Pages 208–215 of: Proc. First Annual Telecommunications Conf.*
- McClellan, S., Burst, Ken, & Grimes, Gary. 1999 (July). Issues and Techniques in Network-Based Distributed Healthcare: Advanced Networking Techniques. *In: Proc. 4th World Conference on Integ. Design & Process Technology*.
- Nichols, Kathleen. 1999 (February). *Using Premium Service to Provide Internet2 QoS*. www.internet2.edu/qos/premserv.html.
- Sikora, John, & Teitelbaum, Ben. 1999 (February). *Differentiated Services for Internet2 (draft)*. www.internet2.edu/qos/diffserv.html.
- Winokur, T., McClellan, S., Kujawski, E., Siegal, G., Listinsky, C., Reddy, V., & McDonald, J. 1996. A prototype telepathology system combining features of dynamic and static telepathology. *Cell Vision: J. Analytic Morphology*, 3(6), 474.