

Quality of Service in Networks and Distributed Systems

Andrew Campbell Domenico Ferrari Stefan Fischer
Lars Wolf

October 2, 2000

Abstract

Distributed multimedia systems are becoming more and more important in many situations of our daily life, for instance in office applications (video conferencing), learning environments (tele-teaching and tele-learning, virtual universities), or entertainment (online games, video-on-demand). Usually, some of the media types used in such an application have specific requirements on their transmission and presentation. The notion of Quality of Service (QoS) plays a central role when discussing about how to fulfil these requirements of multimedia applications. Distributed multimedia systems need QoS support in order to function properly. Moreover, other applications such as certain simulation systems need QoS functionality as well.

For this reason, research in QoS has increased significantly during the past few years. For an end-to-end QoS, which is in most applications necessary (user to user), support has to be provided in all components of the participating systems, i.e., the end system components, the communication system and the application. Accordingly, there has been active QoS research in network hardware (switches, routers), protocol software (RSVP, RTP etc.), operating systems (CPU scheduling), user interfaces, etc. Today, some of the basic technical issues are understood, but a significant amount of work is still necessary. Furthermore, additional research is devoted to (partially) non-technical issues such as pricing for QoS, but also new technical developments such as Active Networks.

This seminar concentrated on these new issues. In 30 talks and a number of panel discussions, the following topics were covered: QoS Architectures/QoS Management, Integrated and Differentiated Services, Multicast and Routing Issues for QoS, QoS in Mobile and Wireless Environments, Pricing and Accounting for QoS, QoS in Heterogeneous Networks, Active/Programmable Networks and QoS, QoS in Middleware, User Level QoS, and Adaptive Applications. This report gives an outline of all talks and thus provides a good overview of the current state of the art in QoS research.

Contents

1	QoS Issues in a Mobile Environment	4
2	Scalable Management of Differentiated Services Networks by using Dynamic Aggregation	4
3	Quality of service management for distributed multimedia applications	5
4	On the Benefits of Using MPLS to Carry Interdomain Traffic	5
5	A Broker Architecture for Value-Added Internet Services	6
6	Charging and Accounting for QoS-enhanced IP Services	6
7	The End-to-End Principle in Network Pricing System Design	7
8	Traffic Shaping in IP Networks	8
9	Charging for Services (with or without Quality)	8
10	Systemic Quality of Service	9
11	Changing Priority of IP Packets and Splitting Traffic	9
12	RSVP Signalling for Commercial QoS	10
13	Design Issues for Edge-Based QoS	10
14	Dynamic Edge Provisioning for Core IP Networks	10
15	Active Segmented Adaptation	11
16	On Formal Methods for QoS-adaptive Systems	11
17	QoS-aware Approaches in Middleware	12
18	Load-tolerant Differentiation with Active Queue Management	12
19	QoS & Heterogeneity – A Phenotype Approach	13
20	Pricing Reservation-based Services	13
21	QoS ad hoc internetworking: dynamic adaptation of differentiated services boundaries	14
22	Effort-limited Fair (ELF) Scheduling for Wireless Networks	14

23 Pricing QoS-Differentiated Services	15
24 L4+ QoS Detection for HTTP and RTP	16
25 Time-Dependent Resource Reservation in QoS-Enabled IP Networks	16
26 TCP's Contribution to Wide-Scale Self-Similarity	16
27 Group Voice Conferencing in Mobile Environments	17
28 An Architecture for Providing QoS in the Future World Wide Web	17
29 End-to-End Delay in Multiplayer Games – or – Are there any general QoS issues in networked games?	18
30 Heterogeneous Services for Group Communication	18

1 QoS Issues in a Mobile Environment

Nicole Beriér, GMD IPSI, Darmstadt, Germany

I presented some QoS issues we work on within a MobileIP reference installation. I discussed the integration of home services with the services of a foreign network with respect to performance aspects. An integration/extension of a RSVP implementation with advanced and third party signalling methods was proposed to support mobile signalling. Adaptive applications to support/specify consistency QoS were discussed.

2 Scalable Management of Differentiated Services Networks by using Dynamic Aggregation

Roland Bless, University of Karlsruhe, Germany

The basic mechanisms for deployment of Differentiated Services (DiffServ) in the Internet are already specified by the DiffServ IETF Working Group, but there are still many open related management issues which have to be addressed before Differentiated Services can become fully operational in a global manner. Specifically, dynamic provisioning of some end-to-end services (e.g., 'Virtual Leased Line' services) requires admission control for resources from end-to-end. The presented approach of an integrated management of DiffServ services within and between DS domains is based on 'DS Domain Managers' as dedicated nodes for this task. In order to achieve an aggregation level for the required signalling between these nodes that corresponds to the aggregation levels in the data paths, a novel concept of 'dynamic aggregation' is introduced. Resource reservations for a particular service are aggregated automatically along common paths (not only sink trees). Because state information about contained flows exists only at aggregation endpoints, there is no signaling with intermediate domains (usually Autonomous Systems - ASs) required. Currently, this approach is evaluated by simulations, especially the chosen heuristics and parameters for aggregation criteria and the additional amount of advance resource reservation for aggregates. Moreover, further evaluation will be carried out by using an actual implementation.

3 Quality of service management for distributed multimedia applications

Gregor von Bochmann, University of Ottawa, Canada

The notion of "quality of service" (QoS) was first introduced to qualify properties of the service provided by communication networks, such as loss rate, delay and jitter. However, it turns out that the "quality" experienced by users of communication services, such as teleconferencing, access to remote multimedia databases or video-on-demand, depends also largely on the performance properties of the end systems (client workstations and servers) that are involved in the application. In order to globally manage the QoS for distributed multimedia applications, it is therefore necessary to consider all the system components involved in the application (the communication network being only one of them). It is also important to consider the user preferences, since different optimization criteria may be used, including the trade-off between quality and cost. The talk will give an overview of the lessons learned from two research projects, one on QoS management for applications involving distributed multimedia databases, and one on multimedia multicasting to large numbers of users (e.g. teleteaching applications). The issues related to the management, at the application level, of the user-perceived qualities are discussed (including cost) and avenues for future research are indicated.

4 On the Benefits of Using MPLS to Carry Interdomain Traffic

Olivier Bonaventure, Université Notre-Dame de la Paix, Namur, Belgium

Multi-Protocol Label Switching (MPLS) was initially proposed as a method to improve the usability of IP over ATM-type networks. Since then, its main utilization has been to provide VPN and traffic engineering services in single ISPs. We show in this talk that MPLS could also play a key role to carry interdomain traffic. We analyze the cost of using MPLS in this environment and show that a single trigger-based approach to flow establishment can be used to control the signalling while still capturing a huge amount of traffic. However, a key issue to be considered is the high burstiness of the interdomain MPLS flows. The burstiness could be a problem when using guaranteed bandwidth MPLS flows.

5 A Broker Architecture for Value-Added Internet Services

Torsten Braun, University of Berne, Switzerland

Virtual Private Networks (VPNs) are spanning over various different locations and can be established over the Internet by IP Security (IPSec) tunnels offering secure, i.e. encrypted and authenticated datagram transport. VPNs have been realized in the past using leased line services providing guaranteed QoS to the users. Therefore, IP-VPNs should also provide similar QoS. Differentiated Services (DiffServ) seems to be an appropriate technology to meet these goals. However, it becomes difficult, if VPNs have to cover different locations which are interconnected via several Internet Service Providers. In addition, users should be able to set up, modify, and release VPNs on demand.

To achieve these goals, a management architecture has been developed which is based on so-called service brokers. Service brokers can be considered as generalized bandwidth brokers which are not only able to negotiate bandwidth but also other services such as secure IPSec tunneling. Service Brokers interact with the users in order to meet their demands and with other service brokers in order to negotiate service level agreements (SLAs) among ISPs. In addition, service brokers have to control network provisioning and charging. The service broker implementation provides several different interfaces, e.g. to web servers or RSVP routers which can signal particular service needs to the service broker.

An important issue is signalling among the service brokers in order to achieve scalability. For example, frequent modifications of a reservation for a single application level flow should not result in similar frequent modifications of SLAs among ISPs. Signalling approaches based on slight overprovisioning, thresholds and traffic measurements have been developed. Simulation results show the suitability of the approach.

6 Charging and Accounting for QoS-enhanced IP Services

Georg Carle, GMD Fokus, Berlin, Germany

The presented work aims to support commercial provisioning of QoS enhanced IP services. New elements for charging and accounting are presented that allow to provide a flexible and scalable charging and accounting Internet infrastructure. A wide range of charging schemes will be used concurrently in the future Internet. The decision of a

provider to select one of the known charging schemes depends on many factors such as service type, user preferences, expected offered load, and business strategy. In order to support a wide range of charging schemes, the accounting infrastructure has to be flexible and adaptable to the different requirements. A framework for policy-based charging and accounting is presented, where the functionality of metering, collecting of metered information, accounting of usage data, applying of charging formulas, and billing can be configured using policies. A Classifier Configuration Notation (CCN) is presented suitable for configuring meters. CCN configuration can be translated automatically into a meter specific configuration language. In order to improve metering performance, a technique is proposed for splitting metering tasks between different functional entities such as user space metering and kernel space filtering. The accounting and metering infrastructure is complemented by an approach for distributing tariff information using the Tariff Formula Language (TFL). This allows users to select services according to a currently valid tariff.

7 The End-to-End Principle in Network Pricing System Design

Jon Crowcroft, University College London, UK

Traditional telecommunication networks implement metering, accounting and charging together with admission and access control at the exchange.

However, in the Internet, routes and provider selection may be dynamic, routes are certainly dynamic, and inter-provider routes are often asymmetric. Intelligence is concentrated at end systems, and is to be avoided at all costs in the core. Based on these observations, we have started a large European project called M3I (the Market Managed Internet), which is building prototypical systems for dynamic pricing more in keeping with the Internet design guidelines.

We describe an architecture for pricing network usage for cost recovery and for user incentive pricing which pushes all functionality such as metering, admission control, and adaptation to the edges of the network, and right into the end systems. By securing the code that "franks" packets, and delivering it to end systems on demand from third party risk brokering agencies, we allow arbitrary policies (for example, different rules for splitting cost between sender and receiver) to be implemented and an evolving market place in the very pricing system itself. This permits a good match with experience in other novel networks such as the mobile telephone market, where a constellation of different contracts has emerged.

We describe how one example of a policy can be implemented with this approach, namely proportional fair pricing. With this, we can implement an efficient and fair

service differentiation. Future work will examine actual user acceptability of congestion pricing in networks, and the viability of risk brokering as a means to offering predictable, stable prices in the evolving service market.

8 Traffic Shaping in IP Networks

Anja Feldmann, University of Saarbruecken, Germany

Managing large IP networks requires an understanding of the current traffic flows, routing policies, and network configuration. Yet, the state-of-the-art for managing IP networks involves manual configuration of each IP router, and traffic engineering based on limited measurements. The networking industry is sorely lacking in software systems that a large Internet Service Provider (ISP) can use to support traffic measurement and network modeling, the underpinnings of effective traffic engineering. In this talk I describe the AT&T Labs NetScope, a unified set of software tools for managing the performance of IP backbone networks. The key idea behind NetScope is to generate global views of the network, on the basis of configuration and usage data associated with the individual network elements. Having created an appropriate global view, we are able to infer and visualize the network-wide implications of local changes in traffic, configuration, and control. Using NetScope, a network provider can experiment with changes in network configuration in a simulated environment, rather than the operational network. In addition, the tool provides a sound framework for additional modules for network optimization and performance debugging. We demonstrate the capabilities of the tool through an example traffic-engineering exercise of locating a heavily-loaded link, identifying which traffic demands flow on the link, and changing the configuration of intra-domain routing to reduce the congestion.

Joint work with Albert Greenberg, Carsten Lund, Nick Reingold and Jennifer Rexford.

9 Charging for Services (with or without Quality)

Domenico Ferrari, University of Piacenza, Italy

We describe the preliminaries of an investigation into the feasibility of designing a flexible client service provider charging interface. The work was motivated by the desire to add such an interface to a system that allowed new services to be added to a pool of existing application services. The general model of charging we came up with is based on the "room" analogy: getting a service is similar to entering a room the door of which is equipped with a clock reading mechanism and a counter; the latter

counts accesses, the former measures the dwelling time of the client inside the room. More complex services may be described by models consisting of several rooms, each inside the other. The general formula of the charge corresponding to each room is $C = hA + kd$, where h, k are coefficients (h is a binary one), A is the access charge, and d is the duration of the client's stay in the room. We examine also the relationship that must exist between the parameters of the charging models used by the transport service provider and by the application provider in order for the latter not to lose money. The work that remains to be done is also reviewed.

10 Systemic Quality of Service

David Hutchison, Lancaster University, Great Britain

There is an increasing need to be able to build large-scale computing systems. These systems are demanded in major application domains including telecommunications, transport, healthcare, banking, and retail. The success of these systems can be measured in terms of performance, dependability, security, scalability, and evolvability. All of these are aspects of QoS as applied to the entire system, i.e., "systemic" QoS.

Much current QoS research is applied to specific characteristics such as performance, and to specific levels such as the network. This presentation discusses the need for systemic QoS, arguing that there is a pressing need for new approaches to QoS in-the-large. New approaches should be able to benefit from the impressive body of research in QoS management, both completed and underway.

Promising avenues for systemic QoS research include the study of system theory and practice, consideration of the work of the dependability community, and learning from the QoS contracts and enterprise modelling work of the Open Distributed Processing (ODP) community.

It is important to reach out and attempt to work with researchers and users from other disciplines such as mathematics.

11 Changing Priority of IP Packets and Splitting Traffic

Jorma Jormakka, Technical University of Helsinki, Finland

Offering several classes of service with different QoS levels in the Internet requires development of new traffic control methods. In the talk, two methods for traffic control, are presented: changing priority queueing which raises the priority and may delay a packet, and correlated traffic splitting. Simulation and analysis of the methods have been described.

12 RSVP Signalling for Commercial QoS

Martin Karsten, Darmstadt University of Technology, Germany

In this talk, I have presented a brief overview of a general signalling architecture, which is based on an extended version of RSVP. Particularly, this architecture allows to integrate a variety of QoS technologies, including DiffServ and IntServ. Additionally, I have explained the capabilities of a new implementation of RSVP, which is able to sustain the signalling load of more than 50,000 flows on standard PC hardware, without any problems. These results in combination with our earlier work in the area of pricing and charging for a multi-service Internet lead to the preliminary conclusion that it is well feasible to realize a multi-service Internet based on this architecture, both technically and economically. However, much more work is needed to gain real-world experience by experimenting with the combination of various QoS technologies.

13 Design Issues for Edge-Based QoS

Edward Knightly, Rice University, Texas, USA

In the last decade, a large body of work has been devoted to providing quality of service to individual real-time flows. However, due to deployability and scalability concerns, architectural support for management of individual flows in the network core is not forthcoming. In this talk, I will discuss design issues for pushing admission control functions to the network edge, describe several possible solutions, and identify key challenges for future work.

14 Dynamic Edge Provisioning for Core IP Networks

Raymond Liao, Columbia University, New York City, USA

Effective edge capacity provisioning is an important architectural component of the emerging Differentiated Service Internet. We propose a set of dynamic provisioning algorithms, which are operational at the edge routers of a differentiated services core network. These edge mechanisms include: i) ingress dynamic link sharing, which augments class based queueing techniques with bandwidth utility functions so that dynamic link sharing can be used to distribute bandwidth among traffic conditioners located at edge routers; and ii) egress dynamic capacity dimensioning, which formulates bandwidth dimensioning at egress links to peering/ transit networks taking into account measured core network traffic conditions. We demonstrate through analysis and simulation that the proposed edge provisioning architecture is efficient and effective at supporting user link sharing policies while taking into account the dynamics of traffic load measured in the core network.

15 Active Segmented Adaptation

Hermann de Meer, University College London, Great Britain

The concept of Application level Active Networking (AAN) has been developed in the ALPINE/ALAN project to introduce flexibility and openness into networking architectures without compromising performance and security of core networking functions. AAN is applied in the presented study to solve the congestion control problem in heterogeneous best-effort and differentiated service networking environments. Segmented adaptation matches the scope of service provisioning and the granularity of service differentiation as opposed to traditional end-to-end and flow-based (TCP) schemes. While TCP congestion control does not similarly reflect particular service assurance levels, TCP-friendliness is nevertheless considered a desirable property of segmented adaptation.

16 On Formal Methods for QoS-adaptive Systems

Jan de Meer, GMD FOKUS, Berlin, Germany

QoS-adaptive systems embody reflexive controlling mechanisms. They are called reflexive because their decision-making components "reflect" the current dynamics

of the system state on which internal controlling decisions are taken. There are two basic concepts of controlling: forward and backward controlling. Forward controlling is oriented along the flow of information, and backward controlling feeds information backwards to the source of the information flow.

17 QoS-aware Approaches in Middleware

Klara Nahrstedt, University of Illinois at Urbana-Champaign

We have presented two different architectures and approaches in middleware. The first one is a reservation-based approach included in the $2k^Q$ middleware architecture. The basic model consists of a *QoSProxy*, which performs (1) QoS translation between the user/application and system QoS requirements, (2) service management with configuration and discovery capabilities, and (3) service coordination of reservations, and *Resource Brokers* such as CPU broker and communication broker to provide for resource reservation. The second approach is an adaptation-based approach included in the Agilos middleware architecture. Agilos is a three-tier architecture. The first level includes adaptors and observers to adapt data based on adaptive control theory. The second level includes configurations which enforce functional adaptation in case of sudden change in resource availability. The third level includes entities such as gateway and negotiations to ensure appropriate distributed control. Both approaches are currently integrated in order to provide a wide QoS support in a heterogeneous computing and communication environment.

18 Load-tolerant Differentiation with Active Queue Management

Olov Schelen, Lulea Technical University, Sweden

Current work in the IETF aims at providing service differentiation on the Internet. One proposal is to provide loss differentiation by assigning levels of drop precedence to IP packets. In this talk, we evaluate the forwarding mechanisms RED In and Out (RIO) and Weighted RED (WRED) in providing levels of drop precedence under different loads. For properly controlled low drop precedence traffic, RIO and WRED can be configured to support sheltering (i.e., low drop precedence traffic is sheltered from losses caused by traffic load at other precedence levels). However, if traffic control fails (e.g., due to inaccuracies in admission control or topology changes) such configurations can cause starvation of traffic at high drop precedence levels. Configuring WRED to instead offer relative differentiation can eliminate the risk of starvation. However, WRED can-

not, without reconfiguration, both provide isolation and avoid starvation at overload of low precedence traffic. To achieve this, we propose a new queuing mechanism, WRED with Thresholds (WRT). The benefit of WRT is that, without reconfiguration, it supports isolation when low drop precedence traffic is properly controlled and relative differentiation otherwise.

19 QoS & Heterogeneity – A Phenotype Approach

Jens Schmitt, Darmstadt University of Technology, Germany

Heterogeneity in large decentralised systems as e.g. the Internet is a matter of fact, though it could also be viewed as design principle for such systems in order to allow for continuous evolution. Network QoS approaches have so far mostly dealt with homogeneous environments and treated heterogeneity with respect to differing network QoS systems as an "accident". In this talk a different view is taken where heterogeneity for network QoS systems is acknowledged as factual and a generic approach to dealing with the issues of interworking network QoS systems is presented. Furthermore, it is argued that this interworking should be regarded at the phenotype level of network QoS systems, i.e. along their properties and not along the mechanisms that are involved (which could be viewed as the genotype of the QoS system). This results in more generic solutions to interworking problems, as is shortly outlined by the example of consolidating the different granularities of QoS systems.

20 Pricing Reservation-based Services

Henning Schulzrinne, Columbia University, USA

In a public network, reserving resources has to incur incremental cost compared to best effort service to allow efficient network resource use and prevent "hogging" of resources. However, traditional resource reservation systems posit a fixed price for service, which tends to either lead to underutilized networks or rejecting calls. We explore making the service price a function of network load. More precisely, our service model assumes that price consists of an access charge, a holding cost reflecting the opportunity cost if a reserved resource remains unused by the reserver and a congestion charge that rises above zero as the network load approaches one. Applications are assumed to be aware of the utility function of the user and thus reduce their bandwidth demands as prices increase. Since many QoS-sensitive applications, such as multimedia applications, cannot change their bandwidth on short time scales without inducing noticeable "fading" artifacts, we assume that prices are constant over periods similar to the typical soft state refresh periods, i.e., about 30 seconds to a few minutes. Prices are set

iteratively, using a tatonnement mechanism, each period based on resource requests. Applications can then reserve resources based on the current price. Thus, we do not require applications to reveal their utility curves to the network and do not require synchronization of requests. Signaling and computational effort is modest and similar to soft state mechanisms such as RSVP. We describe aggregation methods that further reduce the message volume.

We have designed and implemented a resource reservation protocol, RNAP, that allows to experiment with different pricing strategies. The protocol can act as a bandwidth broker protocol in a diff-serv environment or in a more "in-band" fashion.

We also envision that intermediaries may assume the risk of price fluctuations, at a premium, so that end users see constant prices.

We present simulation results that show that the system increases efficiency for both producers and consumers of bandwidth and allows to trade off connection rejection rates and price variation. Bandwidth is shared fairly among applications, depending on the elasticity of the user's demand.

21 QoS ad hoc internetworking: dynamic adaptation of differentiated services boundaries

Michael Smirnov, GMD FOKUS, Berlin, Germany

QoS internetworking is a feature, the Internet needs to become QoS aware in the end-to-end sense. Only one issue of this - adaptation of the differentiated services architecture, in particular, dynamic creation of DiffServ virtual boundaries is addressed below. Our solution is fully distributed and data driven, therefore it could be considered as QoS ad hoc internetworking on contrary to statically configured DiffServ. The proposal is two fold. First, to support invariance under aggregation we propose to maintain Per Domain Behaviours (PDB) based on per path behaviours, and, second, to use group communication based on native IP multicast for needed QoS signalling. The paper shows that due to flexible grouping policies the approach has high scalability and good deployment potential.

22 Effort-limited Fair (ELF) Scheduling for Wireless Networks

Peter Steenkiste, Carnegie-Mellon University, Pittsburgh, USA

While packet scheduling for wired links is a maturing area, scheduling of wireless links is less mature. A fundamental difference between wired and wireless links is that wireless media can exhibit substantial rates of link errors, resulting in significant and unpredictable loss of link capacity. This capacity loss results in a special challenge for wireless schedulers. For example, a Weighted Fair Queue (WFQ) scheduler assumes an error-free link and specifies how flows should share the link capacity. However, this specification is not sufficient to determine the correct outcome when link capacity is sharply reduced, because flows that have been allocated the same weights may differ greatly in their ability to tolerate throughput loss.

In this paper, we first describe the wireless scheduling challenge in terms of an effort-outcome disconnection. Next we propose a novel notion of fairness for wireless links, effort-limited fairness (ELF), which extends WFQ via dynamic weight adjustments. ELF guarantees that all flows experiencing an error rate below a per-flow threshold receive their expected service, defined as a specified rate for reserved flows or a specified share of best-effort capacity for best-effort flows. After motivating and defining ELF, we present a practical approximation algorithm, which we evaluate through both trace-driven simulation and measurement of a prototype wireless radio network based on the WaveLAN physical layer.

This is joint work with David Eckhardt.

23 Pricing QoS-Differentiated Services

Burkhard Stiller, ETH Zurich, TIK, Switzerland

Pricing determines the economic principle for differentiation of services. Based on the change of traffic characteristics of advanced applications (from best-effort to guaranteed services) traffic control is required. However, this control requires the management of state and its sometimes exclusive usage. To provide the right incentive to utilize those resources on valid incentives, pricing and monetary-based mechanisms may avoid users from wasting "public" resources. Based on a model of functional units required for an Internet charging architecture, two extreme pricing modes are presented. One for in-advance determination of volume-dependent charges and its precise calculation and approximative solution based on truncated Markov chains. The other one for highly dynamic auction-based approaches achieving a pareto-efficient result

for multi-class multi-provider scenarios. Both approaches show the clear feasibility in the Internet today, and an efficient performance in technical and economic terms. Concluding, the particular use of these schemes in an open market will be controlled by demand and various business models. Moreover, the presented approaches provide excellent policies for tackling congestion and for providing cost recovery mechanisms for Internet Service providers offering differentiated services in an open market.

24 L4+ QoS Detection for HTTP and RTP

Heinrich Stüttgen, NEC Europe, Heidelberg, Germany

L4+ QoS Detection is a scheme to utilize QoS information contained in higher layer protocols (L4+) to provide Network QoS or CoS support. The scheme has been implemented on a Linux Router and has been integrated with an existing DiffServ router. Using this prototype we can show, that Web traffic (HTTP) and real-time multimedia streams (RTP) can be properly detected, separated and provided with the right, user-defined class of service.

25 Time-Dependent Resource Reservation in QoS-Enabled IP Networks

Giorgio Ventre, University of Napoli, Italy

A number of distributed applications require communication services with Quality of Service (QoS) guarantees. The QoS provisioning issue in the Internet has been addressed by the IETF with the definition of the Integrated Services (IntServ) and Differentiated Services (DiffServ) frameworks. In the actual Internet suite model, resource reservation mechanisms on which these models are based are totally time-unaware. Yet, we believe that, in some cases, associating a time interval to network resource reservations could be useful for both users and network providers. In this talk we present a distributed scheme for time-dependent reservations in QoS-enabled IP networks. We also show how the standard signalling protocol RSVP may support this new reservation style, with only a few minor modifications, so that time could be effectively considered as a QoS parameter. Finally, we present a first prototype implementation of the major component of the proposed architecture and we provide some hints on future applicability scenarios of the advance reservation paradigm and its impact on related topics such as policing and charging techniques in QoS-enabled IP networks.

26 TCP's Contribution to Wide-Scale Self-Similarity

Andras Veres, Columbia University, New York City, USA

In this talk we have presented our results related to how TCP congestion control helps asymptotic self-similarity to be carried from a generation point to distant areas. The main mechanism presented is adaptivity and end-to-end control. We demonstrated that TCP adapts to background processes well approximated by a linear system above a characteristic time scale which depends on the path properties. Consequently, TCP takes on the correlation structure from the background process.

27 Group Voice Conferencing in Mobile Environments

Ulrich Walther, International University in Germany, Bruchsal, Germany

We present the design and the implementation details of a group conferencing tool for the Heidelberg virtual tourist guide. The application uses only 50% of the bandwidth of a 9.6 kbit/s GSM connection thus leaving resources for other applications that need to access network resources concurrently. The system architecture with its conferencing server and the clients is described as well as the QoS issues that were crucial during the design of the application.

28 An Architecture for Providing QoS in the Future World Wide Web

Klaus Wehrle, University of Karlsruhe

The talk discussed the problems arising when quality-based IP services should be integrated or used in the World Wide Web. For example, for signaling the service information from the client to the web server, a simple user interface is needed. Furthermore, the TCP protocol, which is mostly used by the services used in the web, can not fully use the whole reserved bandwidth that the client is paying for.

The presented architecture tries to solve these problems. The service information, for instance, will be signalled within the HTTP GetRequest coded as standard HTML parameters, as with CGI scripts. This avoids defining new signalling protocols and closes the gap between the user and the interfaces to reserve bandwidth. For an easy handling by the client, two approaches have been presented, e.g., a modified web client called

QoSzilla. Thirdly, the problem of poor TCP performance could be solved by adapting the TCP protocol in the web server, which is a simple point of change. By removing the congestion control (which is not useful on guaranteed bandwidth links), adapting the sending window to the path capacity and integrating a rate-based flow control, the TCP protocol could nearly use the whole reserved bandwidth.

29 End-to-End Delay in Multiplayer Games – or – Are there any general QoS issues in networked games?

Lars Wolf, University of Karlsruhe, Germany

Computer-based games are already a big market and there are expectations from various sides that this will increase even more. Multiplayer games where several persons interact simultaneously will receive much interest since competing with human counterparts is typically considered as much more interesting and challenging than playing just against a computer. Network-based multiplayer games, especially via the Internet, will be important. A major problem with respect to networked-based multiplayer games is caused by the network transmission delay. This means that it takes a while until the information, e.g., about the new position of the opponents objects, reaches the receiver. While some technical means can reduce this delay, some delay will always exist. This delay causes several problems and leads to paradoxical situations. Hence, a primary goal is to provide for a global consistent status of the game – independent of the existing delays.

Not only from a technical point of view, the boundary between games and 'serious' applications is not always a sharp one. For example, flight simulators may be seen as games but also as (semi) professional training tools. Hence, we believe that the developed techniques cannot only be applied to games but also to more serious applications.

Current games are not able to handle the requirements of real-time games such as RC car racing simulators sufficiently. Work at INRIA introduced an approach which uses a static synchronization delay. Our investigations have shown that such a simple approach is not sufficient to fulfill the requirements of the human players because then the behavior of the game is unacceptable. We designed an approach with an adaptive synchronization delay which improves the system. Using an additional approach, by using extrapolation, we reached the best results. However, these methods are relatively application specific. This raises questions about more generic solutions, about the dependence on the specifics of setups, how congestion within multiplayer games can be treated, etc.

30 Heterogeneous Services for Group Communication

Martina Zitterbart, Technical University of Braunschweig, Germany

Current communication systems can be characterized as rather inflexible in terms of services that are provided. However, an increasing variety in networked applications leads to a higher differentiation with respect to network services needed. As a result, a higher flexibility of communication systems is needed to properly serve the increasingly diverse service requirements. Especially, new and possibly individual, application-tailored services should be easy to integrate into communication systems without requiring changes in the underlying communication infrastructure. With AMnet, a framework is presented, that provides an open programmable platform allowing the provision of a variety of application-tailored services. AMnet uses programmable as well as active networking technologies. Programmable networking is used for the provision of services, i.e., a network node is programmed according to the applications needs. Programming is done through out-of-band signalling so that regular data transfer is not affected by it. For signalling purposes, active networking is used. Signalling messages are equipped with so-called evaluation programs that determine how suited a programmable node is for provisioning the required service. With this method, it is also possible to locate services that are currently available. Regarding AMnet, a programmable network node has been designed and implemented as prototype. First experiments have been carried out with services, such as QoS filtering for video streams, provision of dedicated and adaptive error control for specific links as well as indirect transport protocols for devices that are attached through a wireless link.