

Report Dagstuhl Seminar 02441

Quality of Service in Networks and Distributed Systems
27.10.2002 – 01.11.2002

Organisatoren: A. Campbell (USA)
S. Fischer (D)
K. Nahrstedt (USA)
L. Wolf (D)

Introduction

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 endsystem 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 report on the Dagstuhl seminar on "Quality of Service in Networks and Distributed Systems" gives an excellent overview on the state of the art in QoS research. It contains abstracts of 23 talks which dealt with most of the above-mentioned research topics. Included were talks on QoS right on the network level, especially in wireless networks, such as those from Stefano Basagni from Northeastern University on QoS in Bluetooth networks or from Jörg Diederich, TU Braunschweig on a simple and scalable handoff prioritization scheme in mobile networks. On the other end of the spectrum, a number of application-oriented approaches were presented, such as the one given by Torsten Braun, University of Berne, on IP Telephony over Differentiated Services or by Ralf Steinmetz, TU Darmstadt, on media semantics. And in between these extremes, many other topics were covered which are worth looking at. We propose to do that and have a thorough look at the report. all authors would certainly be happy to be contacted for further information. And now, enjoy!

Braunschweig, April 23rd, 2003

Lars Wolf Stefan Fischer

Stefano Basagni (joint work with Chiara Petrioli), Northeastern University, Boston, U.S.A.

Multi-Hop Scatternet Formation for Bluetooth Networks

We describe a new algorithm, BlueStars, for the formation of ad hoc networks of Bluetooth (BT) devices. The described solution is fully compatible with the BT specifications and achieves the following desirable properties. The protocol is executed at each node with no prior knowledge of the network topology (distributed and localized). The choice of the masters is driven by the suitability of a node to be 'best fit' for the role of master. The resulting scatternet is connected (if the original topology is) and it is a mesh, thus achieving robustness. In the talk we also evidentiate how much the initial phase of device discovery heavily affects the following phases of piconet formation and piconet interconnection. We conclude with a brief discussion of performance evaluation results.

Rita Wouhaybi (joint with Raymond Liao and Andrew Campbell, Columbia University, New York, USA)

Incentive Engineering in Wireless LAN Based Access Networks

Traffic regulation in public and private wireless LANs face a number of significant challenges, particularly in commercial networks where there is a need for efficient regulation of bursty transactional applications, support for bandwidth reservation services while inhibiting bandwidth hogging by mobile devices, and incentivizing user cooperation. In this research, we take a new approach to solving these problems by applying incentive engineering techniques to wireless access networks. We design two incentive – based allocation service classes: an instantaneous allocation (IA) class, which provides better throughput, and a stable allocation (SA) class, which provides better allocation stability. Our approach possesses a number of beneficial properties including minimizing the algorithmic and protocol overhead on mobile devices, Nash bargaining fairness for the IA service, and incentive compatibility for mobile users promoting the truthfully selection of services class and bandwidth declaration. We use analysis simulation and experimental results from a wireless testbed to demonstrate the effectiveness of wireless incentive engineering.

Jens Schmitt (TU Darmstadt)

Reactive Load Control based on Binary Marking

Existing approaches to achieve Quality of Service (QoS) guarantees in packet-switched network as e.g. the Internet are based on proactive resource allocation mechanisms, in particular admission control, traffic regulation, and packet scheduling. The existing resource allocation model of the Internet, however, is a reactive one based on TCP's flow control mechanisms respectively active queue management mechanisms like random early detection (RED). This approach assumes elastic, i.e. adaptive applications, whereas many applications demanding QoS exhibit inelastic utility functions.

We investigate how a new approach based on traditional reactive mechanisms and so-called load control gateways can deliver QoS to inelastic applications. First experimental results are very encouraging in terms of high scalability and evolvability of such a reactive load control system.

Mark Doll, Uni Karlsruhe (TH)

Web-based initiation of quality-based communication

Assuming a per-flow end-to-end reservation architecture we are trying to deduce requirements for this QoS management architecture that have to be fulfilled to allow for also using the QoS architecture for non QoS-aware applications. We present a web-based approach to provide the glue between, management architecture “user” preferences and legacy applications. It turns out that the main difficulty is the support for ephemeral ports, that is not restricted to some selected legacy applications. The gain is that by modifying only two applications - browser and web server - all applications can profit from an QoS-enabled Network, though presenting a feasible migration strategy from no QoS to full QoS. The management architecture just has to support third party initiated reservations and some sort of transport facility for the initial setup of a trust relationship.

Stefan Fischer, TU Braunschweig

Increasing the Connectivity in Ad-Hoc Networks

Connectivity is one of the major concerns in ad-hoc networks. If for instance a certain link is made available by two adjacent mobile stations, these stations typically don't have any major interest in keeping this connection; they rather have to fulfil their application task. So, often applications relying on that link might quickly loose this connection. We propose, therefore, to introduce stations into the network, whose only job is in the maintenance of the network; i.e. they do not have any application-dependant task. In the talk we present simulations that show the feasibility and effectiveness of the approach. Connectivity can be largely improved. We also describe several ways to improve and influence these "robots" behaviour.

Jörg Diederich, TU Braunschweig

Simple and Scalable Handoff Prioritization in Mobile Networks

A major problem when providing Quality of Service in mobile networks is the occurrence of handoff drops, i.e., there are not sufficient resources in the new cell after a handoff. Handoff prioritization is one possibility to deal with this problem. However, existing approaches are either not scalable with regard to the number of handoffs per session, or they are highly complex to administrate due to too many design parameters to configure. This talk describes the 'Simple and Scalable Handoff Prioritization Scheme' SiS-HoP, which is well-suited for networks with small cells. It is furthermore easy to administrate and can be deployed incrementally. As a result, SiS-HoP can admit more sessions in low-mobility scenarios and reduces the number of handoffs in high-mobility scenarios, compared to handoff prioritization schemes without a mobility prediction.

Kurt Geihs, TU Berlin

Middleware – Quo(s) Vadis?

We look at application-level quality of service and how it can be supported by middleware extensions. Our approach is to build a generic QoS framework for middleware architectures

such as CORBA and .NET Remoting. We provide specification means and QoS implementations which are integrated at various points in the middleware architecture. Furthermore, our work also addresses the platform-independent specification of QoS-enabled application components. A new UML-based specification tool supports this task and the mapping from the platform-independent to the platform specific model as well as the code generation for a specific middleware platform. The approach is evaluated in a commercial application scenario as part of an EC project.

Henning Schulzrinne, Columbia University
(joint work with **H. Tschofenig, X. Fu, R. Hancock, et al, Siemens**)

CASP – Cross-Application Signaling Protocol

We propose a new generic signaling protocol, CASP (Cross-Application Signaling Protocol) suitable for a wide variety of state establishment purposes. We envision CASP as being used for reserving resources on a per-flow or per-class basis, diagnosing network behaviour, depositing active network programs in nodes and managing firewalls and NATs. CASP separates next-node-discovery from signaling, using a “scout” subprotocol. Scouting is only needed if the routing protocol does not indicate the next CASP node.

CASP uses standard reliable signalling protocols such as SCTP or TCP to avoid having to reinvent congestion control, flow control, fast loss recovery and other standard transport protocol features.

An initial specification has been submitted as an Internet Draft.

Jörg Liebeherr, Univ. of Virginia
(joint work with: **A. Burchard, R. Boorstyn, C. Oottamakorn, S. Patek, C. Li**)

Recent Progress on a Statistical Network Calculus

The presentation gave an overview of our efforts in developing a statistical network calculus. The goal of the statistical calculus is to derive expressions for delay and backlog in a network, using the concepts that were developed in the 1990s for a worst case analysis of networks. We discussed our notions of effective envelopes and effective service curves, which are, respectively, upper bounds for aggregates of traffic arrivals and probabilistic lower bounds for service guarantees. Both bounds are probabilistic bounds, but are expressed as non-random functions. We demonstrated that effective envelopes and effective service curves can exploit a substantial amount of statistical multiplexing gain.

Ulrich Fiedler, TIK, ETH-Zürich (CH)

QoS and Transient Simulations of Web Traffic: Using Quantiles to Characterize User-Perceived Latency in Simulations with Heavy-Tailed Input

Simulations with web workloads, which use input generated by sampling a heavy-tailed object size distribution, remain in transient state over all feasible periods of time. This means that statistics that depend on moments of the object size distribution, as the average object size or the average user-perceived latency of downloads, do not converge in any reasonable time. We therefore investigate whether quantiles of user-perceived latency are suitable statistics for the performance evaluation in such simulations. We exploit the fact that quantiles of a heavy-tailed distribution do not depend on the moments of a distribution. We show that

quantiles of interest in samples from a heavy-tailed distribution converge to a normal distribution in reasonable periods of time. Hence, if latency is approximately proportional to the object size, latency quantiles in simulation output also converge to a normal distribution. Therefore, we propose a method to reliably test for this convergence. We validate this method by a simulation study which shows convergence if the network utilization is not too high. Our work suggests that latency quantiles are indeed promising statistics to evaluate simulations with web traffic.

Roland Bless, Universität Karlsruhe (TH) Institut für Telematik

Dynamic Aggregation of Reservations for Internet Services

Providing guaranteed end-to-end-services with QoS on demand is not a trivial task, especially in the Internet. While Differentiated Services may solve the basic problem of scalable QoS provisioning in the data path, management of Diffserv domains and interdomain provisioning of services on demand is still an open problem. Particularly a scalable QoS signaling scheme is required. The DARIS architecture provides a management approach that dynamically creates aggregates of reservations between Autonomous Systems. This reduces the number of reservation states as well as the number of reservation signaling messages. Simulations showed that the savings will be, even in the Internet, in the order of several orders of magnitude. Currently, the signaling protocol (that also provides special support for aggregation) is modified and enhanced for mobile nodes.

Torsten Braun, Universität Bern, CH

IP Telephony over Differentiated Services

This talk described a concept based on adaptive applications and Differentiated Services in order to provide Quality-of-Service required by real time applications. Assuming the availability of different ordered service classes, we propose that a real time application always selects the lowest and cheapest service class that still can meet the application requirements. The selection depends on RTP-based QoS monitoring and on additional probing of the quality of the next lower service class. The service class selection algorithm has been implemented within an IP telephony application and has been evaluated in a test-bed consisting of Linux-PC based Differentiated Services routers. The evaluation proves the suitability of the approach but also allows several issues for further improvement

Klara Nahrstedt, University of Illinois at Urbana Champaign
Joined work with Xiaohui Gu and Jingwen Jiu

QoS- aware Service Composition

Distributed Multimedia Services are becoming an integral part of the ubiquitous computing and peer-to-peer environments. The environments are of dynamic nature consisting of heterogeneous mobile end systems. This characteristic requires a flexible service composition framework that provides end-to-end QoS. We have developed an unified QoS aware service composition framework which consists of two tiers: the composition tier and the distribution tier. The composition tier deploys a service coordination and consistency algorithm to provide QoS consistency between peer services as well as selection of best abstract service path if multiple entities of a service exist. The distribution tier represents the mapping of the

abstract service path into the physical ubiquitous or P2P proxy environment. We have investigated an approximation algorithm for the K-cut mapping between the abstract service graph and K ubiquitous devices. In the PLP environment we have developed a distributed hop-by-hop service path finding algorithm, strong utility load balancing and resource information (computing CPU, memory and bandwidth) for appropriate proxy selection along the service path. Our experimental results in the smart room environment and large PLP simulation confirm that the unified QoS-aware service composition is feasible and delivers desirable QoS guarantees.

Dimitrios Pezaros, Lancaster University, U.K.

dp@comp.lancs.ac.uk

Measurement Monitoring & Control for Improved QoS (On the applicability of in-line measurements for Next-Generation Networks)

The ability to measure monitor and control the service experienced by network traffic, is becoming of increasing importance due to the widespread IP connectivity and hence the increased need for service management, as well as due to the candidacy of the Internet to provide a global telecommunications infrastructure, supporting a range of Different services and traffic, with assured QoS. Current Measurement techniques face difficulties in assessing the actual (real) performance and service offered by the underlying infrastructure either due to their limited, single-point monitoring nature, or due to their liability in producing misleading results, coming from measuring special-purpose (synthetic) traffic. The proposed hybrid approach for conducting network measurements tries to overcome the limitations imposed by current architectures. We theoretically prove and practically trying to evaluate that we can provide for accurate, network service measurements based on the latest, extendable protocol & systems deployment. By defining and using IPv6-based, in-line measurements we aim at revealing the true properties of Internet traffic and stress their particular applicability for next Generation networks as well as for mobile wireless environment, where resource are in shortage.

John Cushnie DMRG, University of Lancaster, UK

(joint work with David Hutchison)

(www.comp.lancs.ac.uk/computing/users/cushnie/)

Provision of QoS for Wireless Access Networks

With the introduction of low-cost IEEE802.11b wireless equipment, wireless access networks are being implemented to allow high bandwidth public and corporate access to the Internet. Most wireless access networks are best-effort only services relying on the high bandwidth connection to the Internet to provide adequate Quality of Service (QoS) to the users. As the number of users and / or the traffic being carried increases the QoS provided by the network is reduced. Traffic control and shaping may be employed within the wireless access networks to provide a level of QoS appropriate to the users budget (financial or otherwise) and the classes of network traffic being carried. This talk presented a lightweight network architecture that may be used to provision QoS to clients in a wireless access network based on the required priority and available bandwidth. The proposed architecture implements monitoring and management of the wireless access network and the connected clients to enable effective traffic control and shaping to be achieved.

Ralf Steinmetz, TU Darmstadt, Multimedia Communications Lab

Towards Semitransparency & towards the Use of Media Semantics for QoS of Media Streams

In our vision 'QoS' will be improved by

1. analysing in detail QoS at the network edges, and at the access network, and taking somehow QoS at the backbone network for given
2. using semantics of data (a description with metadata will be common, see e.g. semantic web) for the 'regulation' of traffic and the adaptation of streams.

In our work on caching, gleaning and adaptation of data streams some results concerning the subjective quality of scalable video were shown. This information is used to control e.g., which data should be resend at all, ahead of resending other data in order to correct the data located at the caches and proxies involved.

Lars Wolf, TU Braunschweig

Much research on QoS has been done over the last years. Various schemes and systems have been developed mainly for QoS support in networks but also for other parts such as middleware or operating systems.

However, target application scenarios and resource capacity situations have changed over time. Looking at wireline networks, demands for QoS and effort for provisioning are in a mismatch, especially considering the capacity available in backbone networks.

Areas which are still useful for QoS research are wireless access networks, and future systems such as sensor networks. However, for application support in wireless networks, per-flow support will be needed in order to support the application well. In opposite to the trend and discussions over the last years, scalability is less of a problem, respectively must be seen in a different light.

Ulf Bodin, Lulea University

Hints and Notifications (HAN)

With current Internet protocols, users may experience low and unpredictable forwarding quality at wireless links. This is due to varying link properties caused by changing radio conditions. Decreasing forwarding quality can cause severe degradation in utilization. Allowing applications and transport layers to communicate with wireless link layers can improve forwarding quality and utilization.. We propose to enable such communication by adding HAN to the Internet architecture. Hints can be implemented through IP options, and notifications through ICMP messages.

Burkhard Stiller

University of Federal Armed Forces Munich, Germany and ETH Zürich, Switzerland

Auditing and Charging in the A^x Architecture

While the existing AAA Approach of the IETF is quite limited with respect to the integration of important components such as auditing or charging, advanced architectures are needed. The A power x Architecture (standing for AAA and beyond, with x a placeholder for auditing, charging, and privacy) proposes a generic policy-based approach. The talk details major components and illustrates the need due to the high demand on mobile services and remote access to Internet services. In particular, charging is discussed in detail. Furthermore, the IST Moby Dick project provides some implementation platform for studying those principles. This work was joint work with my Ph.D. student Hasan, ETH Zurich, and an academic guest, Christoph Rensing, from TU Darmstadt.

Anthony D. Joseph
University of California, Berkeley

In the Griffin project, we are exploring the development of a network analysis and emulation that provides developers with a highly accurate emulation of a wired or wireless link. Griffin also provides applications with predictive information about network conditions and topologies that enables them to adapt BEFORE conditions deteriorate or improve. Preliminary results show that Griffin techniques can help applications reduce wide-area bandwidth requirements by up to 75%, improve performance on slow networks by up to 90%, and significantly improve performance over lossy networks.

The Griffin architecture includes components of several levels of the network stack, from link and media access layers through overlay networks, and up to applications. Of particular interest is Tapas, Griffin's network monitoring, modelling, emulation, and prediction system. Using Tapas provides two significant benefits to protocol and application developers by enabling the development of new protocols using highly accurate simulation techniques and simplifying the development of predictive adaptive applications.

Michael Smirnov , Fraunhofer FOKUS

Policy Group Control Issues

How can we define semantics & interfaces between components of Internet composite (more than a single pair of client-server involved) service? Consider each component as a combination of component's mechanism and a set of policies - rules defining a choice in the behaviour of a mechanism. Separating concerns further, we distinguish between behaviour (obligation) policies and safeguard (authorisation) policies, triggered by events produced under control of notification policies. How should we design for evolvability of such a system? A scenario based service design is proposed : a service is implemented scenario-wise as a collection of behaviour and notification policies, and a set of safeguard policies, achieving high component re-use along with high levels of system openness & flexibility. The approach is implemented as GEN (Group Event Networking) platform; demonstrated with three programmable examples: service creation; load balancing; and distributed computation of a configuration policies derived from an SLA.

Geoff Coulson, Lancaster University

A Middleware-based Approach to Programmable Networking

NETKIT is a new project that is investigating the application of component technology to programmable networking environments. In terms of motivation, we observe a number of deficiencies in the current state of the art in software platforms for programmable networking. For example: i) existing approaches to programmable networking are often partial (e.g. classical active networking does not support distributed signaling or coordination protocols; similarly, open signaling systems often do not support downloadable code), ii) existing approaches do not adequately support manageability (e.g. they don't have generic support for configuration, instantiation, reconfiguration, adaptation, extension, evolution, removal of programmable networking software), and iii) existing approaches are typically language specific (e.g. C++ frameworks for in-band packet handling, or Java for AN execution environments). To address these deficiencies (and others that will be discussed in the talk) we propose a middleware-based approach to programmable networking. In more detail, we are exploring the consistent use of a middleware-derived component model to build functionality at all levels of the programmable networking architecture, from basic OS support on routers, to in-band packet handling, to execution environments to (middleware-supported) signaling and coordination protocols. The component model supports the interoperation of components written in any language and also provides generic services for (remotely) managing programmable networking software in all the above-mentioned areas. To be realistically efficient (especially in demanding areas like in-band packet handling), and to support fine-grained components, the component model must also add negligible overhead. In this talk I describe the NETKIT approach in detail and discuss progress to date and upcoming challenges.

Juergen Quittek, Network Laboratories, NEC Europe Ltd.

Next Steps in QoS Signalling: Do we need RSVP version 2?

The presentation classifies and briefly evaluates the Resource reSerVation Protocol (RSVP) and discusses some shortcomings including complexity, scalability and lack of openness. This is followed by a report on the IETF activity called Next Steps In Signaling (NSIS). The NSIS working group has developed requirements for a successor of RSVP. The requirements were derived from application scenarios, such as terminal and session mobility and 3G wireless networks, that were not in the focus when RSVP was developed. The requirements include a modular, extensible design, that is independent of any QoS signaling paradigm and they include higher security. Finally, feasibility and performance studies of a protocol meeting these requirements are presented. They were conducted by implementing and measuring operation of a simple protocol meeting most of the new requirements. Still the IETF RSVP working group needs to decide if it considers it necessary to start working on a new version of RSVP.