

Building Open Programmable Multimedia Networks

Andrew T. Campbell, Aurel A. Lazar, Henning Schulzrinne and Rolf Stadler
The COMET Group
Center for Telecommunications Research
Columbia University, New York, NY 10027-6699
comet.columbia.edu

Abstract

Recent advances in distributed systems and transportable software and increasing demand for better quality-of-service (QOS) control in multiservice networks are driving a re-examination of network software architectures. We established the COMET Group (Control Management and Telemedia) at Columbia University's Center for Telecommunications Research to provide a comprehensive understanding of network software architecture of the 1990s and beyond. In this paper, we present an overview of our activities focusing on new research initiatives, international forum participation and on-going research projects. Collectively, these activities shape our vision of a new era driven by open programmable networking.

1 Introduction

The COMET Group [1] of the Center for Telecommunications Research at Columbia University was established in response to the need for a comprehensive understanding of the software architecture of the networks of the 90s and beyond. Led by Aurel A. Lazar, the group encompasses more than thirty researchers including Henning Schulzrinne and Andrew T. Campbell, Rolf Stadler, visiting researchers, Ph.D., M.S. and B.S. students.

Recent advances in distributed systems and transportable software together with increasing demand for better control of QOS in multiservice networks are driving a re-examination of network software architectures. A new opportunity exists to reconcile the perspectives of the computing and communication communities in new network architectures that support service creation, QOS control, and the joint allocation of computing and communications resources.

Future multiservice (multimedia) networks will carry traffic such as video, audio and computer data with a diverse range of *quality-of-service (QOS)* requirements. In addition, they are expected to have a very large number of physical and logical entities, services and users to be managed and controlled. We believe that the development of open programmable multimedia networks is one of the key research challenges that faces the networking community as we move towards the new millennium. To address this chal-

lence the COMET Group has initiated a number of new projects and international forums to promote the ideas which drive our research. In what follows we describe some of these new initiatives, forums and research projects.

2 New Research Initiatives

The COMET Group has played an important role in establishing new international forums to promote and address the need for research into open architectures, open signalling and network programming for ATM, Internet and mobile networks. The group's work on open programmable networking has been instrumental in the formation of the new IEEE standardization initiative on "Programmable Interfaces for Networks" [8].

2.1 Broadband Networking

The group's research in ATM-based broadband networks has considerably expanded beyond investigating network architectures, and network control and management, into the area of open signalling and telemedia. The work in open signalling was prompted by the simple observation that the origin of the prevalent thinking in signalling and service creation dates back to the late sixties and is based on the fundamental assumption that the "intelligence" required for service creation only resides inside the network. We are convinced that this assumption is now outdated and that a fresh approach to the subject is required.

In response to this challenge the COMET Group played an important role in establishing an international working group called *OPENSIG (open signalling)* [5] for actively exploring network programmability and next generation open signalling technology for service creation and deployment on ATM-, Internet- and mobile multimedia networking platforms. In that spirit the COMET Group aims to make the process of service creation, deployment, control and management of new network services as easy as building and launching new applications on a PC. The OPENSIG group meets on an informal basis every six months to exchange ideas and discuss appropriate forms of collaboration with the stated goal of promoting research and development on open signalling and service creation.

As a result of the tremendous success of OPENSIG, the IEEE Communications Society is sponsoring a new conference on *Open Architectures and Network Programming (OPENARCH '98)* [6]. The first OPENARCH conference will be held on April 3-4, 1998 in San Francisco and will be co-located and organized in conjunction with INFOCOM'98. OPENARCH will offer a forum for the communication of experimental as well as theoretical results aimed at a better understanding of the overall networking architecture and its realization in software. It will encourage a shift of the processes of service creation, resource allocation and control from ad-hoc solutions to a discipline of network programming.

2.2 Internetworking

The Real-Time Transport Protocol (RTP), developed by the Audio/Video Transport (AVT) working group within the IETF with the help of H. Schulzrinne, is gaining wide-

spread acceptance as the common transport protocol for the delivery of continuous-media information over the Internet, including Internet conferencing, telephony and media-on-demand. Its adoption by the ITU-T as part of its H.323, a first for an Internet standard, ensures its use for Internet telephony and LAN-based conferencing in general. RTP describes a data packet format as well as a control protocol (RTCP) that allows participants to gauge the group size, to establish the identity of other group members and measure the quality-of-service received by members. The rate of the control protocol is self-regulated to a small fraction of the data bandwidth.

RTP is currently a Proposed Standard, soon to be “upgraded” to a Draft Standard. One of the problems to be addressed during this transition is the use of RTP in very large multicast groups, as might be found for CATV-like distribution of audio and video programming. In that scenario, with potential audiences of thousands to several million, the self-regulating feedback mechanism and group size estimation of RTCP ceases to function properly. Also, if a large number of users join a multicast group simultaneously, a large surge of control traffic could overwhelm parts of the network. We have developed and analyzed a number of remedies for these problems. These algorithms ensure that the rate of feedback scales regardless of group size and join timing as well as allow quick audience size estimation. They have been submitted to the AVT group for consideration and will likely be integrated into the next release of the standard.

While RTP provides the common base for data transport, the Internet currently does not have a well-developed set of application-layer signaling protocols for Internet conferencing, telephony and media-on-demand services. We have contributed to two protocols to address both the initiation of conferences and the control of streaming audio and video sessions. Both protocols are loosely based upon HTTP, allowing to reuse both protocol features such as encryption, authentication, billing negotiation and content labeling as well as their server implementations.

2.3 Mobile Networking

Recently there has been considerable interest in the emergence of wireless ATM (WATM) as a possible candidate technology for next generation mobile multimedia communication systems. As multimedia applications migrate to mobile devices, wireless extensions to broadband networks are required to support the seamless delivery of voice, video and data with high quality. In this context WATM is intended to be a direct extension of the existing ATM network with uniformity of end-to-end QOS guarantees.

In early 1996 the ATM Forum established the WATM working group to address the challenge of providing broadband services over an air-interface. The primary goal of the WATM working group is to develop a set of specifications for location management, ATM mobility and a new 25 Mbps air-inteface.

The COMET Group has been actively involved in shaping the discussion within the WATM working group and in the broader mobile and wireless networking community at large. We have participated in the WATM working group since its creation. Our involve-

ment in this new initiative is to track the development of new standards, to exploit the emergence of new techniques through the development of software solutions and, to feed our research results into the working group. Prior to the establishment of the ATM Forum WATM working group we organized the First Workshop on Wireless ATM Networking. The workshop was a great success in bringing together leading researchers involved in the implementation of prototype WATM networks.

3 Software Distribution

The investigations of the COMET Group have been regularly published as papers in leading journals and conferences. A complete list is published in *The COMET Annual Report* that is available on the WWW [7]. A major goal of our research is to provide public domain communication software for broadband, Internet and mobile networks. Our software prototypes also available through the WWW. Currently the following software can be downloaded:

- **kStack**: User Space Native-Mode ATM Transport Layer with QOS Support (<http://comet.columbia.edu/software/kStack/>);
- **RTP**: A Transport Protocol for Real-Time Applications (<http://www.cs.columbia.edu/~hgs/rtp/>);
- **TREX**: Telecommunications Research Exchange: A Dynamic Auctioning Game for Spectrum Allocation in Wireless Networks (<http://argo.ctr.columbia.edu:1024/~nemo/Trex/trex.html>);
- **xbind**: A broadband kernel for devising the next generation technology of service creation, deployment and management on ATM based broadband networks (<http://comet.columbia.edu/xbind/>); and
- **mobiware**: QOS-aware middleware for mobile multimedia communications (<http://comet.columbia.edu/wireless/software/mobiware>).

4 Research Focus

In this section we highlight a number of new COMET research initiatives in ATM, Internet and mobile multimedia networks. Following this we provide an overview of each on-going project in the form of a bullet report. For full details on these COMET research projects see [7].

Broad access to open multimedia networks will play an important role in enabling the emergence of new broadband networking technologies and programming paradigms. A major thrust of the COMET group is the development of an *open programming environment* that facilitates the easy creation of distributed multimedia services on top of ATM, Internet and mobile networks as illustrated in Figure 1. By open, we mean that the proposed platform must support a set of APIs for resource control and management which third party service providers can use when developing new services. By programmable, we mean that these APIs should be 'high-level' enough to allow the service specification and creation process to be carried out via a high-level programming language utilising

distributed systems technology such as CORBA.

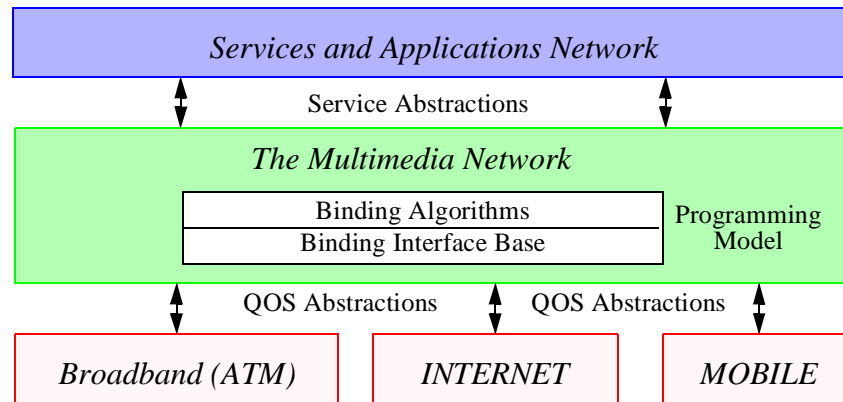


Fig.1 Model of programming for multimedia networks.

4.1 Open Programmable Networking

Since the Fall of 1994 we have been experimenting with a broadband kernel called **xbind**. The broadband kernel is as a programmable operating platform that supports the creation, deployment and management of networked multimedia services (e.g., connection management, route management, admission control, QOS mapping, etc.) and mechanisms for efficient resource allocation. The term ‘kernel’ is deliberately used to draw a parallel between its role as a resource allocator and extended machine, and that of a typical operating system. The broadband kernel behaves as a resource allocator because it mediates and arbitrates between conflicting requests for resources made by various parties in the system. It functions like an extended machine because it provides a simplified means of accessing fundamental system services by abstracting away the operational complexities of provisioning these services.

Resource allocation in **xbind** is based on game theoretic models. By taking an economic (market based and game theoretic) approach in the engineering of multimedia networks, we seek solutions where the intelligence and decision making is distributed and thus scalable, and the objective of a more efficient and fair utilization of shared resources results from the induced market dynamics. Thus, we are borrowing, and adapting the tools of game theory from Economics and using them to solve problems of resource allocation (such as admission control, routing and flow control) in multiservice networks. By viewing a network as a collection of resources which users are selfishly competing for, this approach gives rise to efficient, decentralized algorithms, and leads to network architectures which provide explicit QOS guarantees.

The players in the network economy are software agents, rather than humans. Agents acquire resources, such as bandwidth and buffer space, from the network on behalf of applications (video, voice, data transfer). Under appropriate rules of interaction, the collective actions of all the agents constitute a distributed intelligence, superior to that of any single controller. Thus the challenges are to analyze noncooperative behavior and algorithmic strategies, and to design the mechanisms (rules of the games) that will ensure

the desired outcomes.

4.2 Advanced Internet Services

From modest experiments transmitting Internet Engineering Task Force meetings in 1992, the Internet multicast backbone and the associated set of IP multicast tools have seen widespread use. Research into packet audio and video in the Internet, however, go back to the early 1970s. Based on these experiences, the Real-Time Transport Protocol (RTP) was developed, which is now on the Internet standards track. A number of research organizations and companies have developed interoperable, RTP-based audio and video tools. Also, as an Internet first, the ITU has adopted RTP as a part of their multimedia conferencing architecture for LANs, commonly known as H.323. A number of well-known vendors have started shipping H.323-based Internet telephony products.

The demand for Internet multimedia teleconferencing grows dramatically by the day, but much work needs to be done before truly scalable transport and signalling protocols, and teleconferencing tools are in place. The COMET Group is actively meeting these challenges [3]. Future work in the group includes the development of a *Session Initiation Protocol (SIP)* that performs all standard telephony signalling functions, including automatic and manual call forwarding. In addition, new feedback modes for RTCP need to be developed to address some of the scaling problems that manifest as the number of participants in an RTP session scales from hundreds to potentially thousands. Other multimedia Internet services research includes the development of a *Real-Time Stream control Protocol (RTSP)* for controlling the retrieval of continuous media over the network.

Currently, we are developing a suite of protocols and application building blocks which address these open research issues. Our goal is to provide scalable solutions for current and future multimedia services as the demand for a multimedia-ready Internet accelerates.

4.3 Wireless Media Systems

The COMET group has initiated a new research project in Wireless Media Systems [4] to investigate the provision of QOS in mobile multimedia communication systems. Next generation wireless and mobile communication systems will be required to support the seamless delivery of voice, video and data with high quality. Delivering hard *Quality of Service (QOS)* guarantees in the wireless domain is complex due to large-scale mobility requirements, limited radio resources and fluctuating network conditions. To address this challenge we are developing a QOS-aware middleware platform called *mobiware* which contains the complexity of supporting multimedia applications operating over wireline and wireless ATM networks.

Mobiware a software middleware platform that seamlessly runs on mobile devices, base stations and mobile-capable ATM switches. It is based on CORBA and Java technology that is designed to operate between the application and radio-link layers of future wireless media systems. The platform incorporates new architecture and novel adaptive algorithms to support QOS controlled mobility. Mobiware provides value-added QOS

support by allowing mobile multimedia applications to operate transparently during handoff and periods of persistent QOS fluctuation. The goal of the mobiware adaptive algorithms is to transport scalable flows, reduce handoff dropping and improve wireless resource utilization. We use the term “controlled QOS” to distinguish it from hard QOS guarantees offered by fixed ATM networks. Implicit in the term is the notion that flows can be represented and transported as multi-layer scalable flows at the mobile device. Adaptive algorithms help scale flows during handoff based on the available bandwidth and an application-specific flow adaptation policy. This policy characterizes each audio and video flow as having a minimum QOS layer and a number of enhancements.

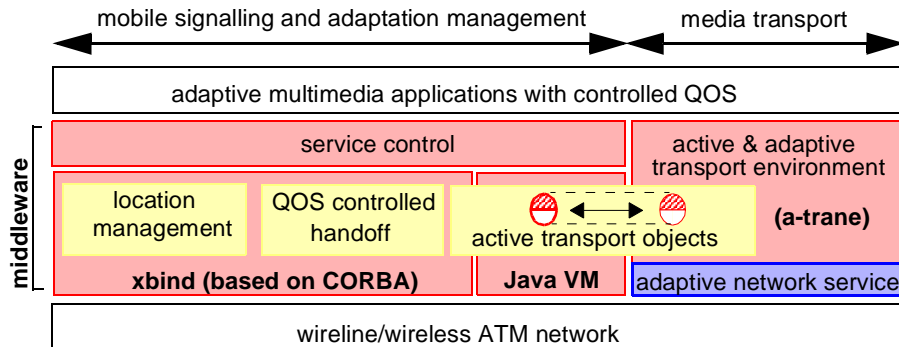


Fig.2 Mobiware

Mobiware promotes the separation between mobile signalling and adaptation management on the one hand and media transport on the other. As illustrated in Figure 2 mobiware utilizes **xbind** and Java for signalling and adaptation management during handoff or periods of persistent QOS fluctuation. The Java Virtual Machine executes on mobile devices, base stations and mobile-capable ATM switches and supports the dynamic execution of *active transport objects*. These transport objects constitute an ‘active’ component of the mobiware transport system called *a-trane* which can dispatch active transport objects to strategic points in the network (e.g., the base station) or end-systems to provide value-added QOS support.

4.4 Prototyping Multimedia Networks

In order to study multimedia networks under various scenarios, such as different load patterns, network sizes and management operations we have built a platform for developing and evaluating control systems for emerging multimedia networks. The platform allows us to closely approximate the functional and dynamic behavior of network control systems and provides support for real-time visualization and interactive emulation. The current implementation runs on the IBM SP2 parallel processor at the Cornell Theory Center, that is connected to a graphics workstation in our laboratory at Columbia University via an ATM link. The platform is currently used in various projects in the COMET group that are all aimed at developing and evaluating network architectures. The approach taken is real-time control based and as such differs from the current work of the Telecommunications Information Networking Architecture Consortium (TINA-C)

which is essentially management based with a focus on data and functional modeling.

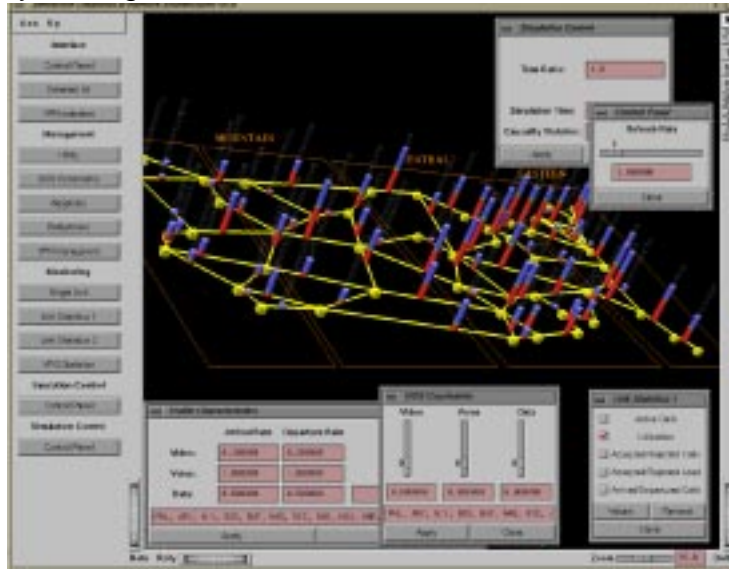


Fig.3 Prototyping Multimedia Networks

The emulation platform consists of four building blocks: parallel simulation kernel, emulation support, real-time visualization and interactive control, and emulated system. The emulated system (as illustrated in Figure 3) and the emulation support modules consist of a set of objects that communicate by exchanging messages, using functions provided by the simulation kernel. The simulation kernel controls the execution of these objects and ensures that messages are processed in the correct order. In order to support real-time visualization and interactive control of the emulated system, the kernel controls the progression of the simulation time, constraining it by the progression of the processor time. The module for real-time visualization and interactive control contains a graphical interface which provides 3D visual abstractions of the system state.

5 COMET Research Projects

The COMET Group is investigating issues arising within the network architecture, network control, network management, resource allocation and networking games, telemedia and information transport, as the basis for building open QOS programmable systems on top of ATM, Internet and mobile networks. In the following section we present an overview of on-going projects in each area. For full details on COMET research see [7].

5.1 Network Architectures

5.1.1 Broadband Kernel

We have built a prototype broadband kernel, *i.e.*, an operating platform that supports the creation of a set of services (*e.g.*, connection manager, route manager, admission controller, QOS mapper, *etc.*). The architecture of the broadband kernel is based on the G-model of the XRM as illustrated in Figure 4. The broadband kernel incorporates the Binding Architecture consisting of an organized collection of interfaces, called the *Binding*

Interface Base (BIB), and a set of algorithms that run on top of these. It also incorporates a native ATM stack for information transport.

BIB interfaces provide an open and uniform access to abstractions that model the ‘local’ states of networking resources. Binding algorithms play a key role in service creation through the process of interconnecting (binding) of networking resources. QOS is explicitly modeled in the architecture via a set of abstractions that characterize the multiplexing capacity of networking and multimedia resources under QOS requirements. The resource allocation mechanisms that are integral part of the broadband kernel will not be discussed in this paper.

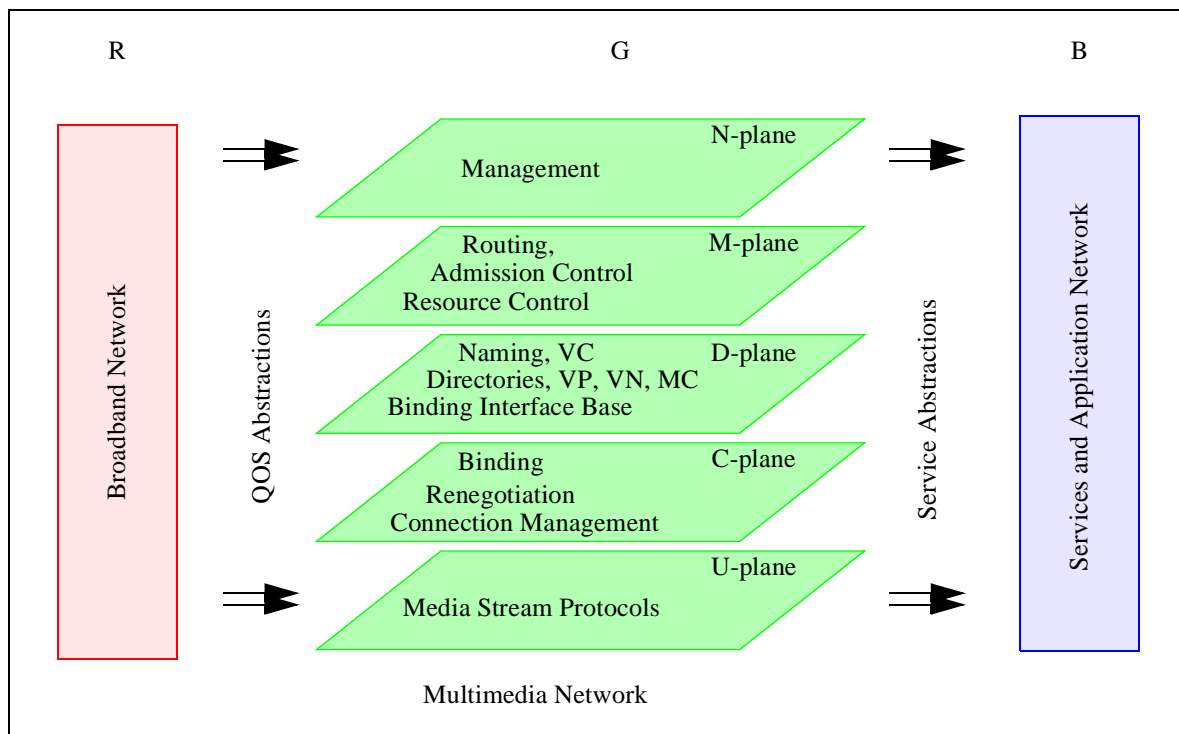


Fig.4 The XRM Model Underlying xbind

We proposed a service creation methodology and demonstrated how scalable multimedia network services can be constructed from a set of broadband kernel services. As a proof of concept, we built a multi-party teleconferencing service upon the broadband kernel. The service creation framework has the following characteristics: binding of different classes of transport protocols to the application, dynamic renegotiation of application QOS using the signalling network and, support for multiple resource reservation algorithms.

5.1.2 Broadband Virtual Private Networks

The goal of this research is to develop a communication architecture that allows the extension of a private high-speed backbone into islands of private component networks, using a public broadband infrastructure. The management and control system of the

resulting global network is independent of the characteristics of the underlying public network services, guarantees QOS including survivability on an end-to-end basis, and exhibits advanced techniques for managing and controlling the interconnections of component networks. The results of this effort are directly applicable to the creation of a virtual enterprise network for a geographically distributed large corporation.

5.1.3 Real-Time Emulation and Visualization of Networks

A major challenge in broadband networking today is to develop architectures that allow for creating services in a flexible way, providing QOS to network applications, and enabling efficient use of network resources. In this project we are investigating interactive emulation and visualization of networks using a supercomputer. Our approach to evaluating and refining architectural concepts is based on experimentation on an interactive prototyping platform. This allows us to study the functionality and dynamics of a network architecture, by experimenting with different network configurations, load patterns, control policies, etc. The experimental platform constitutes an emulator that runs on a parallel machine at the Cornell Theory Center (CTC) and a manager station at Columbia University. It is implemented on an SGI Indigo2 workstation running a 3D user interface for real-time visualization and interactive control. Both machines communicate through NYNET, an ATM network that connects several research laboratories in New York State.

5.1.4 Mobeware

The realization of QOS controlled mobility in WATM networks requires the exploitation of the scalable flow concept and the provision of wireless adaptive services that can best understand the transport needs of scalable flows in a mobile ATM environment. To address these QOS challenges we are developing *QoS-Aware Mobile Middleware* called *mobeware* which takes end-to-end programmability for QOS controlled ATM mobility as its primary design goal. We use the term controlled QOS to distinguish it from hard QOS guarantees offered by fixed ATM networks. Implicit in the term is the notion that mobile flows can be represented as multi-resolution scalable flows. In order to support ATM mobility with controlled QOS *mobeware* provides a set of co-operating QOS-aware transport, mobility and media scaling algorithms. These algorithms operate jointly to provide the seamless delivery of multimedia with QOS control and include:

- *QOS controlled handoff*, provides mobile and ATM signalling, control and management algorithms capable of supporting QOS controlled handoff based on the seamless delivery of media with smooth change in the perceptual quality;
- *media scaling*, exploits the intrinsic scalable properties of multi-resolution audio and video flows and the knowledge of user supplied scaling preferences to filter flows at critical nodes in the network (e.g., base station to mobile) in order to best utilize the available bandwidth;
- *QOS adaptive transport system*, provides a QOS-based API and suitable QOS algorithm for the rate control, flow control, playout control, flow monitoring, QOS manager, error control and filtering for the efficient and flexible transport of audio, video and data to bandwidth restricted mobile devices;
- *adaptive service algorithm*, which best matches adaptive application behavior by

providing hard guarantees for minimum acceptable quality and uses residual capacity to deliver enhanced quality when resources become available; and

- *QoS programmability*, which utilizes distributed systems technology based on the **xbind** platform, implemented by members of the COMET Group, for service creation, resource management and control of wired ATM networks.

Mobiware promotes the division between signalling, control and management on the one hand, and digital audio and video transport on the other. Mobiware is based on extending **xbind** to the mobile environment by providing end-to-end QoS support via resource binding between devices: ATM switches, base stations, multimedia devices (e.g., cameras, display devices), and mobile and fixed devices.

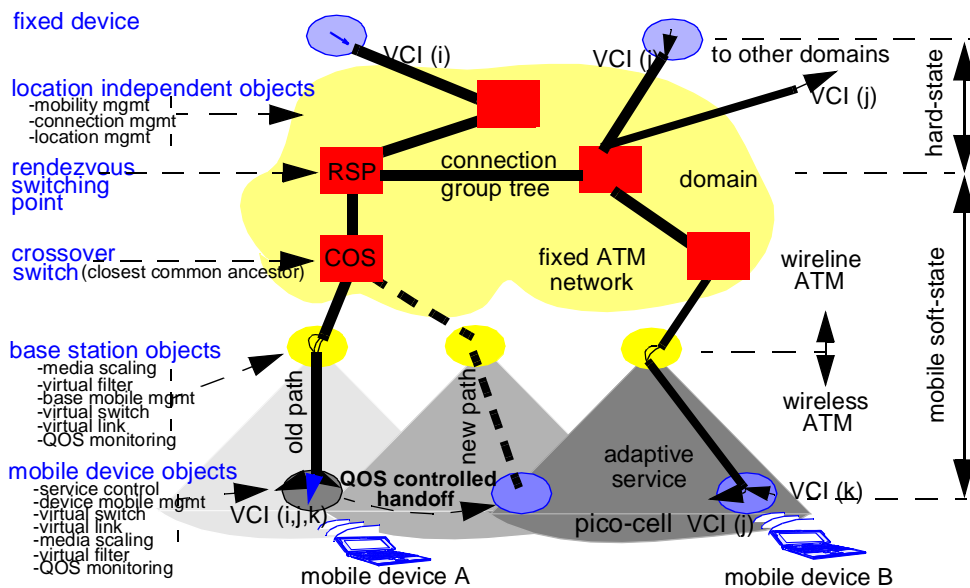


Fig.5 mobiware: QoS-Aware Mobile Middleware

The mobiware platform models the wireless portion of the ATM network as being divided into pico-cells each served by a base station connected to a wired ATM network as illustrated in Figure 5. Base stations are cell relays which translate the ATM cell headers from radio ATM format to that used by standard ATM. Each base station supports signalling, QoS control and management of flows based on semantics of an adaptive service. The existing wired ATM network provides connectivity between base stations

5.2 Information Transport and Signalling

5.2.1 Scaling RTP

Unlike protocols for other conferencing systems, RTP was designed from the beginning to support both large and small groups. Indeed, the scalability of IP multicast and RTP is one of the primary differences between traditional “telecom” conferencing architectures and the Internet model. RTP has a built-in feedback mechanism, with which receivers report on the network conditions. These reports are sent via multicast to all session members. The report interval scales linearly with the media bandwidth, so that reports

consume a small, fixed fraction (typically 5%) of the media bandwidth. This works well for sessions with at most a few hundred participants, even if the media and, thus, RTCP bandwidth is limited to modem speeds. As it happens, most current Mbone RTP sessions fall within this size regiment. However, there is growing interest of using the Internet not just for niche programs such as complete NASA shuttle mission coverage or a university seminar, but also as a replacement for radio stations or cable televisions, with audiences measured potentially in the millions.

Once sessions grow beyond a few thousand, particularly for low-bandwidth sessions, it takes hours to find out how many are actually listening. The long interval between QOS reports from a source makes it difficult to estimate whether, say, the reported packet loss rate is caused by steady, unnoticeable loss or intermittent, disruptive loss events. More disruptively, the current RTCP scaling algorithm leads to a spike in traffic when a large number of receivers join at the same time. We are developing and implementing algorithms that prevent traffic bursts, allow near-instantaneous audience size estimates and new modes of QOS feedback.

5.2.2 Conference Control Tools

Unlike other conferencing system, the Mbone and RTP tools are single-media tools that are used as building blocks for applications. The same audio tool, for example, can be used as an Internet telephone, a tool for Internet group collaboration, an Internet radio receiver and a receiver for media-on-demand, just by changing the surrounding control programs. However, the control of tools is currently rather limited. We are currently building a set of tools that communicate via a conference control bus implemented by host multicast, where information generated by one tool can easily be used by any number of other tools. This makes it easy, for example, to build a shared conference roster application that shows all participants and their media, rather than having each media tool display its own set of participants. Other conference accessories, like talk timers and floor controllers, can be integrated without having to modify the remainder of the applications.

5.2.3 Session Initiation Protocol (SIP)

Today's Internet multimedia conferences are announced in a global multicast directory. There is no way to invite a particular individual to a conference or a phone call, short of sending the person email. Together with Mark Handley of ISI, we have developed an Internet session initiation protocol (SIP) that allows to perform all the standard telephony signaling functions, including automatic and manual call forwarding, call transfer and automatic call distribution, yet is much simpler than the existing telephony signaling protocols. With this protocol, traditional PBXs are no longer necessary where most employees have access to a workstation, PC or network computer, as all call handling is done by the end systems or a simple call gateway at the corporate perimeter. This approach is also much simpler than the computer-telephony integration (CTI) efforts that maintain a separate network for voice. Since the traditional public switched telephony network is going to be with us for many years, interoperation between the two is being investigated. Issues such as gateway selection, billing, security and addressing are necessary ingredients for smooth interoperation.

5.2.4 Real-Time Stream control Protocol (RTSP)

As mentioned earlier, the network protocols and applications can be the same for all continuous Internet media, whether stored as media-on-demand or interactive as telephony and conferencing. Media-on-demand requires additional protocols to retrieve media descriptions and control the flow of media data, for example, to pause or fast forward. We are actively involved in defining the Real-Time Stream control Protocol (RTSP) currently being standardized within the IETF. One of the interesting properties of the protocol is that it re-uses many of the facilities offered by HTTP and thus can share most of the optimized code found in web servers. It also can offer most of the functionality of the MPEG DSM-CC media control suite, yet be described in less than one tenth the page count. In the past, conferencing and media-on-demand were strictly separate activities. RTSP supports the integration of media-on-demand sources into on-going conferences. We are exploring the control and management aspects of this integration.

5.2.5 qstack

The goal of this research is to augment **xbind** with a native ATM transport service. The transport and **xbind** platforms are independent but work in unison to provide end-to-end guarantees for application-level QOS. The transport system includes novel QOS mechanisms for QOS monitoring and adaptation for the delivery of unicast and multicast flows with COMET traffic class QOS. A QOS-based API is provided that shields the application programmer from the complexity of transport QOS management and control.

QOS mapping is needed between different levels of the network in order to provide end-to-end service guarantees. For example, in order to reserve the appropriate network resources at connection establishment time, a connection manager needs to map application QOS requirements into network level QOS requirements. Furthermore, each transport protocol needs to monitor the TPDU level QOS it delivers to the applications to ensure that it provides the appropriate QOS. Thus QOS mapping between the application QOS and transport QOS is also needed.

QOS mapping needs to be done in two phases: a mapping between the type (or category) of services provided in each level and a rescaling of QOS parameters. Currently, we are considering application services such as conferencing (multicast), broadcasting, shared data (e.g., whiteboard). These applications require media streams such as motion JPEG video, MPEG-I video, MPEG-II video, telephony, CD quality audio and reliable point-to-multipoint data transmission. Each application has its own QOS specification: frame rate, frame size, tolerable delay, video frame or audio packet loss and gap loss. The current set of mapping rules is far from complete and consists simply of parameter rescaling proportional to the PDU ratio of each level considered. We are performing experimental QOS measurements concurrently at the application, transport and network level, and are expecting to formulate a set of QOS mapping rules for our set of multimedia applications. We believe that the rules formulated above are the worst case bounds and that experimental measurements will permit us to derive empirical rules that are closer to reality.

5.3 Connection Management and Binding

5.3.1 Binding Architecture

We proposed a conceptual framework (called the Binding Model) for the creation, deployment and management of multimedia services on ATM-based broadband networks with end-to-end QOS guarantees. The model allows the binding of networking resources with the goal of creating distributed services across heterogeneous networking platforms.

The Binding Architecture is a formal description of the Binding Model. Its principal aim is

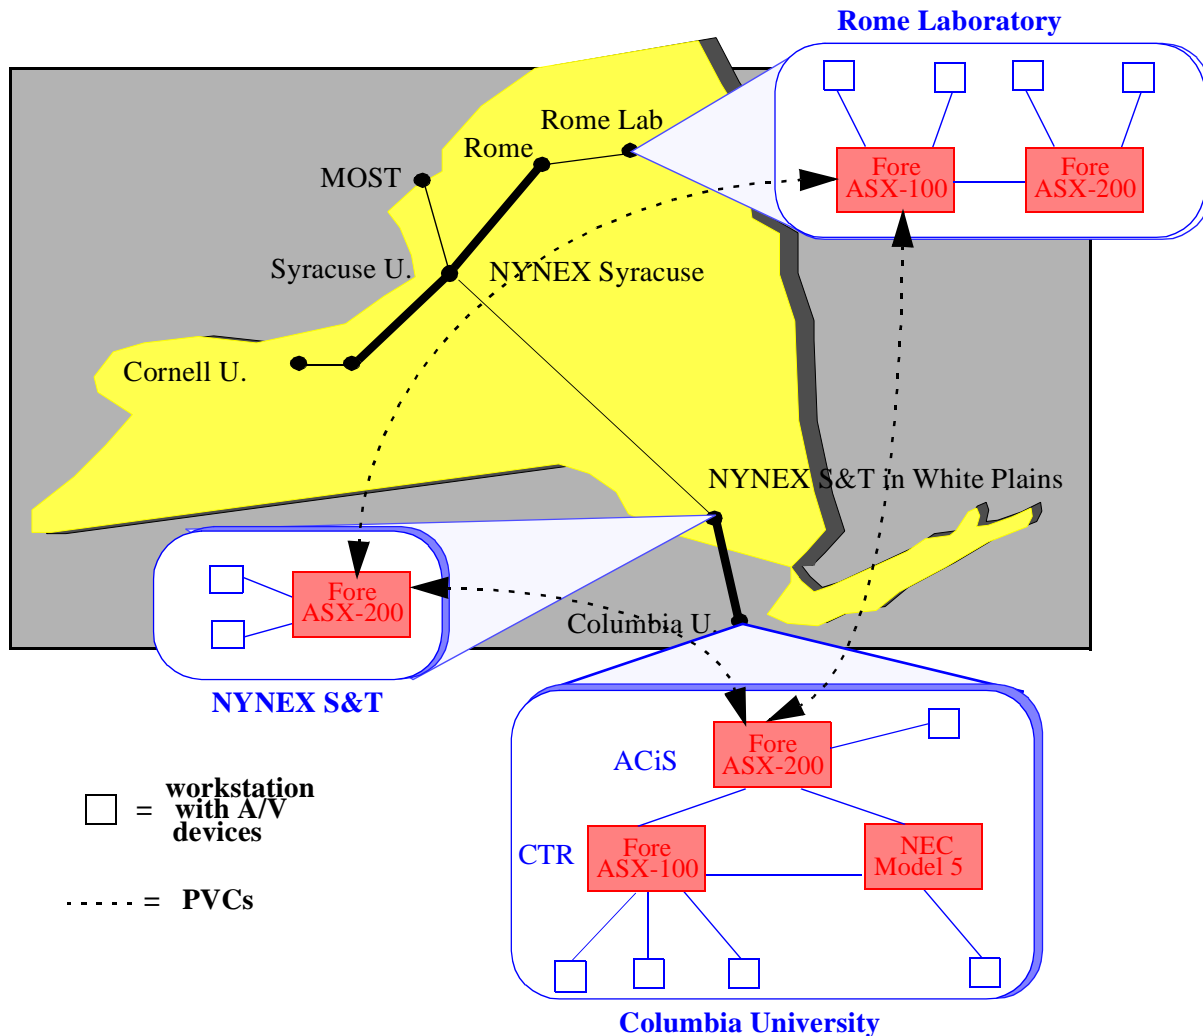


Fig.6 Deployment of the Binding Architecture on NYNET

to provide an open programming environment that facilitates the easy creation of distributed services. It consists of an organized collection of interfaces, called the *Binding Interface Base* (BIB), and a set of algorithms that run on top of these. BIB interfaces provide an open and uniform access to abstractions that model the 'local' states of networking resources. Binding algorithms play a key role in the service creation process through the

process of interconnecting (binding) networking resources. QOS is explicitly modeled in the architecture via a set of abstractions that characterize the multiplexing capacity of networking and multimedia resources under QOS requirements.

Binding Architecture was implemented as part of a middleware toolkit for building scalable multimedia services on top of heterogeneous computing and networking platforms. The Binding Architecture achieves interoperability between the different resource control and management algorithms of a multimedia network by giving these open access to the BIB. By adopting the Common Object Request Broker Architecture (CORBA), an industry standard for distributed computing platforms, the Binding Architecture achieves interoperability in terms of data representation across diverse operating systems and machine architectures.

The first implementation of Binding Architecture was completed at the beginning of 1995 and its binaries were made available in the public domain in the of summer 1995. Starting with the summer of 1995 we have been deploying Binding Architecture on the Columbia Campus ATM network by first installing it on three ATM switches. Since then we have been experimenting with Binding Architecture on the New York State ATM testbed, NYNet, in collaboration with NYNEX S&T and with the Air Force Rome Laboratory. Figure 6 shows the network topology used in some of the experimentation.

5.3.2 Architecting the Control Infrastructure of Multimedia Networks

It is generally perceived that although CORBA based architectures offer much advantages in terms of programming flexibility, they are generally regarded as overly heavy, perceived to suffer from poor performance and thought to be impractical for use in real time telecommunication services. We show that this need not be true and present results of various performance benchmarks that we have conducted for our system as an example. Specifically, we compare the performance of its connection management scheme with published results obtained for UNI 3.0 and demonstrate that careful tuning allows sustained comparable performance.

5.3.3 Multicast

The goal of this work is to design a multicast service architecture that integrates the activities of service access, operation and management on ATM-based multimedia networks. A multicast service provides to one or more users the ability to send information flows simultaneously to more than one receiver. In addition, a multicast service is characterized by a group of senders and receivers, the information flows being sent and the underlying infrastructure needed to deliver a multicast flow to all the receivers under receiver-specific QOS constraints. Three major “components” of our multicast service are group management, flow management and connection management. Collectively these form the basis for multicast session management.

5.4 Resource Allocation and Networking Games

5.4.1 Architecting Noncooperative Networks

Control decisions in large-scale networks are often made by each user independently according to its individual performance objectives. Such networks are called *noncooper-*

ative, and game theory provides the systematic framework to study their behaviour. The operating points of a noncooperative network are the Nash equilibria of the control game. The goal of this project is two fold: (i) to characterize the structure of the network operating points, and (ii) to introduce a unified methodology for architecting noncooperative networks so that their operating points exhibit certain desirable properties. Network control games are typically non-classical constrained games, due to the presence of stability and/or quality of service constraints. Using flow control, routing and admission control as paradigms, existence of equilibria has been investigated for a class of networking games.




Fig.7 Open Spectrum Access Java Applet

5.4.2 Pricing

The problem of *open spectrum allocation* is a challenging area of research that lies at the intersection of network control and telecommunications policy making. As already mentioned, the COMET Group is taking an economic (i.e., market and game theoretic) approach to this problem where the decision making on bidding for spectrum is both distributed and accessible. An *open spectrum access tool* has been developed based on Java (Figure 7 illustrates our open spectrum allocation applet) and deployed on the Internet for experimentation and evaluation of the approach. The tool presents participants with intelligent control algorithms to build manual and automatic bidding strategies. Our intention is to evaluate the open spectrum access system against available auctioning and licensing schemes

5.4.3 Broadband Networking with QOS Constraints

The key objective of this work is to bridge the gap between video traffic modeling and recently developed queueing analysis techniques. Guided by the empirical observation that real-time MPEG video streams exhibit both multiple time scale and subexponential characteristics, we constructed a video model that captures both of these characteristics and is amenable to queueing analysis. We investigated two fundamental approaches for extracting the model parameters: using sample path and second order statistics based methods. Our approach provides a unified context for explaining a number of seemingly conflicting results that have recently been reported in the literature. Along the same lines, our model applies to the full range of queueing behavior that is relevant to engineering different video applications such as real-time video conferencing, video on demand, broadcasting, etc.

5.4.4 Evaluating the Schedulable Region in Real-Time

The schedulable region is the foundation for our network control architecture. It represents the capacity of a multiplexer as the space of possible combinations of calls, and forms an extension of the concept of Shannon capacity. In the present project, we concentrate on evaluating the size of the schedulable region in real-time. Due to the randomness of the calls traversing any given link, the schedulable region of that link may vary over time. We are designing algorithms, which will be run on switches, for measuring the activity in the output ports of switches. This information will be used to determine the size of the schedulable regions of the output links.

5.5 Network and Service Management

5.5.1 Managing Real-Time Services on Multimedia Networks

This project aims at providing high-level control abstractions to network operators. We argue that high-level controls and dynamic visual abstractions are needed for network operators to effectively pursue management objectives in a high-speed multi-service environment. We have developed a management architecture to support these capabilities, based on the concept of managing real-time traffic by tuning the resource controllers in the traffic control system. We have evaluated these concepts on an emulation platform, which includes a management station implemented on a graphics computer and a high-performance parallel machine running the control and management systems of a large network in real-time. This platform allows the experimentation with control abstractions for network operators and the monitoring of management operations in real-time. The next phase of this project will be to port the developed management algorithms to the **xbind** platform.

5.5.2 Building Manageable Multimedia Network Services

The aim of this project is to design a new service management model in the context of Enterprise ATM Networks in which telecommunication services are deployed and managed. We consider four types of services in our design in the first instance: virtual connections, virtual paths, virtual networks, and multicast services. In this context the service management activity is split between affecting the control and state of the system. The initial focus of the work is on configuration and performance management of existing service instances. In this regard the service management activity becomes similar in context to the QOS management activity from a provider viewpoint. The proposed service model is simplified to handle a limited set of entities interacting in a hierarchical structure. The management Application Programming Interfaces (APIs) are defined as object interfaces based on the Common Object Request Broker Architecture (CORBA) distributed platform.

5.5.3 Managing Virtual Circuit and Virtual Path Services with QOS

In this project we investigated an architecture for performance management on wide area ATM networks. The architecture is composed of a signalling system for connection establishment at the Network-to-Network Interfaces that support the creation, deletion and bandwidth control of Virtual Paths (VP). Admission controllers are installed at the

source node of every VP and admit calls that can be guaranteed QOS. A *network management system* based on the OSI standards was used to capture performance information both at the call and the cell level. This information is then used by management system to evaluate network-wide performance measures and apply the appropriate VP controls. An automatic algorithm is capable of providing the necessary VP distribution such that QOS constraints both at the call and the cell level are guaranteed for the observed and forecasted traffic demand. The architecture has been completely implemented on AT&T's Xunet gigabit testbed and experimental results demonstrated.

5.6 Telemedia

5.6.1 Virtual Workshop

The Virtual Workshop (VW) service is a distributed multimedia application that provides various forms of multimedia associations among VW participants. Like a "real workshop" a Virtual Workshop features interactions such as sessions and tutorials, as well as one-to-one conversations among participants. In contrast to a "real workshop" however, the VW participants are physically distributed instead of meeting in the same physical location. They make use of the broadband communication infrastructure and the VW service in order to interact.

The Virtual Workshop is an advanced multimedia service, that requires interaction and cooperation among various entities: the Virtual Workshop application, the network control system, the service management system, and the network management system -- and thus readily allows us to experiment with our interface designs under demanding service requirements. We use **xbind** as a platform for implementing generic management and control interfaces. The **xbind** system already provides us with the fundamental underlying support services necessary in a multimedia network with QOS guarantees.

5.6.2 Service Modelling

The distributed nature of multimedia networks, and therefore of the BIB, suggests many possible means of distributed interactions among different binding algorithms during the service creation process. We have devised a service creation process that consists of five steps:

- The creation of a service skeleton for an application (*e.g.*, virtual circuit, virtual path, virtual network or multicast). For example, the skeleton for a virtual circuit consists of a graph from a source node to a destination node.
- The mapping of the skeleton into the appropriate name and resource space, thereby creating a *network application*.
- The association (or binding) to the application of a media transport protocol, thereby creating a *transport application*.
- The binding of the transport application to resources, creating a *network service*.
- Finally, the binding of the service management system to the network service, thereby creating a *managed service*.

This simple view of the service creation process is surprisingly complete as illustrated in Figure 8.

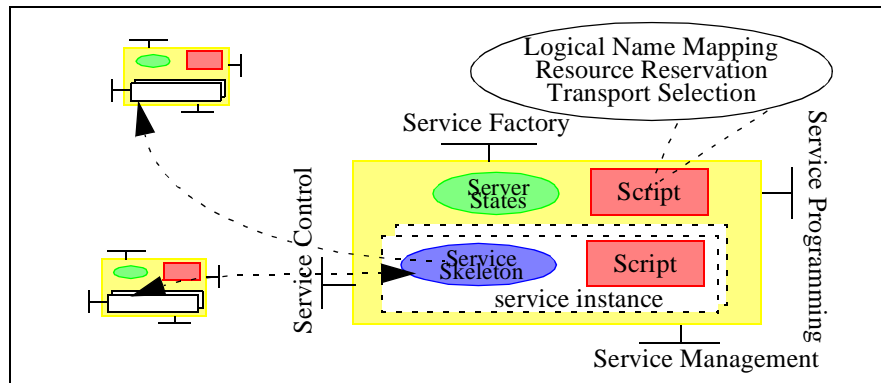


Fig.8 Interfaces for Modeling Servicest

Broadband kernel services and indeed all services are offered by servers which support the creation and maintenance of their state. The interfaces of a server reflect upon the roles that it must play in the process of provisioning a service. Typically these include creation, control, management and programming. Thus, we defined 4 general interfaces called the *Service Factory* interface, the *Service Programming* interface, the *Service Control* interface and the *Service Management* interface respectively. The service factory interface is used to request the creation of a service instance. The service programming interface allows customization or modifications to be made to the algorithmic component of the server or service instance. The service control interface is the operational interface to the service instance and allows the monitoring and manipulation of service instance states during execution. Finally the service management interface allows for monitoring and control of the server and the setting of management policies.

6 Summary

In this report we have briefly described research activities of the COMET group of the Center for Telecommunications Research at Columbia University. All publications, software and information associated with the COMET group is directly accessible via the W3 [1]. In the early 1990s the COMET group developed the necessary theoretical foundations for engineering QOS in broadband networks. Today, we are developing open programmable ATM, Internet and mobile networks that will enable new models for network signalling, control and management to be created in the future.

References

- [1] The COMET Group: <http://comet.ctr.columbia.edu/>
- [2] Open Programmable Multimedia Networking: <http://comet.ctr.columbia.edu/xbind/>.
- [3] Multimedia Internet Services Research: <http://www.cs.columbia.edu/~hgs/research/>
- [4] Wireless Media Systems Research: <http://comet.ctr.columbia.edu/wireless>.
- [5] OPENSIG: <http://comet.ctr.columbia.edu/opensig/>
- [6] OPENARCH: <http://comet.ctr.columbia.edu/openarch/>
- [7] The COMET Activity Report, 1994-1996: <http://comet.ctr.columbia.edu/publications.html>

- [8] IEEE Standardization Initiative on "Programmable Interfaces for Networks", comet.ctr.columbia.edu/opensig/standards/standards.html

Andrew T. Campbell (comet.columbia.edu/~campbell/) is an Assistant Professor in the Department of Electrical Engineering and member of the COMET Group at the Center for Telecommunications Research, Columbia University, New York. He is currently leading a research effort in Wireless Media Systems focusing on the development of QOS programmable middleware for mobile multimedia networks that comprise ad-hoc, broadband and next generation Internet technologies. Before joining academia Dr. Campbell spent 10 years in industry focusing on the design and implementation of network operating systems and communication protocols for packet-switched, local area and tactical wireless networks.

Aurel A. Lazar (<http://comet.columbia.edu/~aurel/>) (S'77-M'80-SM'90-F'93) is a professor of Electrical Engineering at Columbia University. His research interests span both theoretical and experimental studies of telecommunication networks and multimedia systems. The theoretical research he conducted during the 1980s pertains to the modeling, analysis and control of broadband networks. He formulated optimal flow and admission control problems and, by building upon the theory of point processes, derived control laws for Markovian queueing network models in single control as well as game theoretic settings.

He was the chief architect of two experimental networks, generically called MAGNET. This work introduced traffic classes with explicit quality of service constraints to broadband switching and led to the concepts of schedulable, admissible load and contract regions in real-time control of broadband networks. In the early 1990s his research efforts shifted to the foundations of the control, management and telemedia architecture of future multimedia networks. His involvement with gigabit networking research led to the first fully operational service management system on ATM based broadband networks. The system was implemented on top of AT&T's XUNET III gigabit platform spanning the continental US. His management and control research pioneered the application of virtual reality to the management of ATM-based broadband networks. His current research in broadband networking with quality of service guarantees focuses on modeling of video streams and analysing their multiplexing behavior, with emphasis on multiple time scales and subexponentiality.

Professor Lazar is also leading investigations into multimedia networking architectures supporting interoperable exchange mechanisms for interactive and on demand multimedia applications with quality of service requirements. The main focus of this work is on building a broadband kernel (also known as `\bold xbind`) that enables the rapid creation, deployment and management of multimedia services. Resource allocation and networking games algorithms form the core of this architecture. The `xbind` work led to the establishment of the IEEE standards working group on Programming Interfaces for Networks.

Professor Lazar was instrumental in establishing the OPENSIG (<http://comet.columbia.edu/opensig>) international working group with the goal of exploring network programmability and next generation signalling technology. He was the Program Chair of the Fall and Spring OPENSIG'96 workshops and of the First IEEE Conference on Open Architec-

tures and Network Programming (OPENARCH'98,<http://comet.columbia.edu/openarch>).

Henning Schulzrinne (<http://cs.columbia.edu/~hgs>) is an Associate Professor in Department of Computer Science and Electrical Engineering. Dr. Schulzrinne's research interest comprise computer networks, networked multimedia, performance evaluation and distributed systems. Internet protocols and services; congestion control and routing in high-speed computer networks; implementations of multi-media and real-time networks; operating system support for high-bandwidth services with real-time constraints; distributed real-time systems; performance analysis of computer networks and systems.

Rolf Stadler (<http://comet.columbia.edu/~stadler/>) is a research scientist with the Center for Telecommunications Research (CTR) at Columbia University and a principal investigator in the COMET group. He is also an adjunct associate professor with the Department of Electrical Engineering at Columbia University. Dr. Stadler's research interests include multimedia networking and telecom systems; service and distributed systems management; and software architectures of distributed multimedia systems

