

OPEN SIGNALING FOR ATM, INTERNET AND MOBILE NETWORKS (OPENSIG'98)

University of Toronto, Ontario, October 5-6, 1998
<http://comet.columbia.edu/opensig/activities/opensig98.html>

Andrew T. Campbell, Columbia University, USA
Irene Katzela, University of Toronto, Canada
Kazuho Miki, Hitachi Limited, Japan*
John Vicente, Intel Corporation, USA*

ABSTRACT

The ability to rapidly create and deploy new transport, control and management architectures in response to new service demands is a key factor driving the programmable networking community. Competition between service providers may hinge on the speed at which one provider can respond to new market demands over another. The notion of open programmable networks is having broad impact on service providers and vendors across a range of telecommunication sectors calling for major advances in open network control architecture, network programmability and distributed systems technology. In this paper we discuss the origins of the Open Signalling Working Group (OPENSIG) and present a summary of the fifth Workshop on Open Signaling for ATM, Internet and Mobile Networks (OPENSIG'98), which was held at the University of Toronto, Ontario, October, 1998. We also discuss a number of new initiatives in the area of open programmable networks that have recently emerged.

INTRODUCTION

In 1995, a series of workshops entitled Open Signalling (OPENSIG) was started with the goal of making ATM, Internet and mobile networks more open, extensible and programmable. The ultimate goal of introducing open programmable technology into the network is to make the network as programmable as the PC thereby enabling new network services to be launched with the ultimate ease of clicking and downloading an application. The separation of communications hardware from control software is fundamental to reaching this goal; that is, the

separation of the transmission hardware (i.e., switching fabrics, routing engines) from the network programming environment. Such a separation is difficult to realize today. The reason for this is that switches and routers are vertically integrated – something akin to mainframes of the 70's. Typically, service providers do not have access to the switch/router control environments (e.g., IOS), algorithms (e.g., routing protocols) or states (e.g., routing tables, flows states). This makes the deployment of new network services and environments (which may be many orders of magnitude more powerful, flexible and programmable than the proprietary control systems) impossible due to the closed nature of the network nodes.

By modeling communication hardware via a set of open programmable network interfaces, access to switches and routers can be provided thereby enabling third party software providers to enter the market for telecommunications software. By 'opening up' the switches, routers and base stations in this manner the development of strikingly new services and network architectures can be realized. Opening the control of networks through a set of well defined network programming interfaces and distributed programming environments (e.g., CORBA, DCOM, Java) is an approach being developed by the OPENSIG community. In this case, physical network devices can be as abstract as distributed computing objects (e.g., *virtual switches* [xbind,96], *switchlets* [Tempest,96], *virtualBaseStations* [Mobiware,98] and *routelets*

* Visiting Researchers, COMET Group, Center of Telecommunications Research, Columbia University, New York, NY 10027.

[Genesis,98]) using well-defined open programming interfaces. These interfaces allow service providers to manipulate the states of the network objects using middleware in order to construct and manage new network services (e.g., routing, mobility management) with quality of service support.

THE OPENSIG MOVEMENT

Until recently our ideas on signalling and service creation were rooted in the late sixties. They were based on the fundamental assumption that the “intelligence” required for service creation resided inside the network. There was a need to approach the problem of signalling and service creation using new network models and service creation paradigms. Having intelligence only inside the network appeared to be very restrictive.

With that in mind, Aurel A. Lazar called for a working group meeting in October, 1995, to review one of the emerging open signaling paradigms called xbind [xbind,96] being developed by the COMET Group at Columbia University. This review was in part prompted by the interest expressed in this approach by some of the participants of the Columbia Workshop on Telecoms entitled: Binding Architectures for Building Networking Middleware that took place the day earlier. The purpose of the working group meeting was to also discuss appropriate forms of collaboration with the stated goal of promoting research and development on Open Signalling and Service Creation. For information on the agenda discussed and the working group program of Monday, October 23, 1995 see [OPENSIG,95]. During the meeting it was decided to form a working group called OPENSIG whose aim was to do research towards understanding open network control issues as they arise in signalling, middleware and service creation in ATM, Internet and mobile networks.

In October 1995, the OPENSIG working group established an international workshop with the Program Committee of the first meeting Chuck Kalmanek (AT&T Research), Aurel A. Lazar, (Columbia University) and Ian Leslie (University of Cambridge). The workshops took place every six months during 1996 and 1997 (Spring 1996, Fall 1996, Spring 1997 and Fall 1997) in New York, Cambridge and Toronto.

The participation at OPENSIG workshops has included vendors, providers and members of leading academic and industrial laboratories and numbered between 100-150 people meeting informally for two days.

Following the tremendous interest in the OPENSIG activities, the IEEE Communications Society sponsored a new conference entitled: Open Architectures and Network Programming (OPENARCH) [OPENARCH, 97]. The first in the series, OPENARCH'98 took place on April 3-4, 1998, in San Francisco. It was organized in conjunction with and was co-located with INFOCOM'98. Future OPENSIG workshops will take place in the Fall of every year, whereas OPENARCH conferences meets with INFOCOM in the Spring of the same year. The work of the OPENSIG group has focused on the definition, implementation and experimentation of open programmable networks. Having an impact on products and services of the next generation networks is a key goal of the OPENSIG working group. An important step in this quest was the creation of IEEE P1520 Proposed IEEE Standard for Application Programming Interfaces for Networks [P1520,98a].

Recently, OPENSIG members have spun off new technology companies that are commercializing ideas and technologies introduced at the workshops over the last three years. For example, Xbind Inc [Xbind,98] is a high-tech startup located in Silicon Alley, New York and Cplane Inc [Cplane,98] a Bay Area start up. In addition, members of working groups have released source code distributions to the public domain, e.g., software for open programmable broadband [Xbind,96] and mobile networks [Mobiware,98].

OPENSIG'98

The OPENSIG'98 workshop was hosted by the Communications Group (University of Toronto) and Chaired by Irene Katzela. The workshop brought together some new research communities for the first time including members DARPA Active Networks [DARPA, 96] [ANmeets,96] and IETF Differential Services Working Group [Diff,98]. Highlights of the workshop included: a keynote address by Nelu Mihai, (AT&T West Labs) on GeoPlex; reports on new developments in active

networking organized by Ken Calvert (University of Kentucky), an interactive session on views of differential services arranged by Kathleen Nichols (Cisco) and a session on programmable virtual networking organized by Andrew T. Campbell (Columbia University). Other sessions included programmable networking (Chair: Gisli Hjalmtysson, AT&T Research), mobile agents and open signaling (Roger Imprey, National Research Council), open slot session (Chair: Nikos Anerousis, AT&T Labs) and a panel on Industry's view of Open Signaling (Chair: Al. Leon-Garcia, University of Toronto). In what follows, we present a summary of OPENSIG'98.

KEYNOTE ADDRESS

Chair: Aurel A. Lazar, Xbind Inc. and Columbia U.

GeoPlex: An Open Service Platform

Nelu Mihai, AT&T West Labs

Nelu Mihai presented a framework for enabling network-based services called *GeoPlex* which promotes the idea of a *programmable control plane* for Internet. The speaker posed the question of what the future network would look like – a smarter, more active and intelligent network? It was clear that at least the network must keep pace with advancement of end-user services and the requirements they pose on networks. In addition, networks must evolve to rapidly introduce new services - moving beyond a 'bit-transport paradigm'. Nelu Mihai mentioned security problems, complexity, scaling issues and lack of a service creation environment as drawbacks for today's networks. These issues were attributed to the current computing solution paradigm and the inequality of intelligent end-systems and non-intelligent networks. Emerging networks should support middleware architecture providing common capabilities and services where transport carrier technologies are "smarter". In addition, programming interfaces must be built into the network infrastructure and exposed to distributed programming environments.

The keynote speaker described the GeoPlex system as an IP-based service platform that offered support for network (PTT, Internet, Private Enterprise Networks) and media services integration. The GeoPlex system is capable of offering ISPs accelerated development of on-line services. Its service architecture would offer an

advanced Internet services platform supporting an integrated set of subscriber management functions, security, manageability, application and service-level support. The GeoPlex architecture, consists of gateway services that include access control and protocol mediation; core services that provide access, authentication and membership; store services for maintenance and provisioning services; and hops which represented transfer points. Nelu Mihai illustrated how the GeoPlex environment exposes the subscriber management system and service provider management through APIs to IP applications. Concluding, our keynote speaker noted that the GeoPlex system is capable of 'absorbing network complexity', integrating legacy systems and enabling service providers to rapidly deploy and offer new services for the Internet.

ENABLING VIRTUAL NETWORKING

Chair: Andrew T. Campbell, Columbia University

What is a VPN ?, Paul Ferguson, Cisco

Paul Ferguson started the session with a twenty minute tutorial *on Virtual Private Networks (VPNs)* [Ferguson,98]. Providing an industry perspective, he explained the various motivations behind VPNs including the economic, privacy, security and service quality benefits using 'protocol stack' taxonomy to illustrate VPN implementation options. The network layer VPN support included network address translation, GRE, PPTP and L2TP tunnels, controlled route leaking. The link layer VPN support included ATM and Frame Relay virtual connections and link layer encryption. Application encryption and transport-level security protocols were also discussed briefly. The speaker provided a practical perspective for the requirements of VPNs. In summary, VPNs would continue to be a controversial area with alternative implementations, addressing diverse customer needs for privacy, security and quality of service.

Genesis Project: Programmable Virtual Networking, John Vicente, Intel and Columbia University

The second talk [Vicente,98], given by John Vicente, broadened the discussion by introducing the notion of "programmable" virtual networks based on the concept of spawning. A study of active network and

programmable network projects was presented using a taxonomy and generalized framework to put these contributions into perspective for the audience. By way of a survey of these projects, he noted that trends in programmable networks are leading to network infrastructure abstractions and virtualization through open programmable interfaces. The ultimate challenge and abstraction for programmable networks is, in fact, virtual networks. John Vicente spent the second part of his talk outlining a new COMET Group project that he was involved in at Columbia University called Genesis, which he said literally meant “the origin or mode of formation or generation of a thing”. The project was concerned with the development of a *virtual network kernel* capable of profiling, spawning and managing virtual network architectures. Genesis included the notion of parent/child virtual networks (creating nested virtual networks), autonomous control and resources inheritance. Life cycle services profile, spawn and manage virtual networks. While virtual network infrastructure is constructed from routelets interconnected by virtual links.

The NetScript Approach to Programmable Networks, Danilo Florissi, Columbia University

Next, Danilo Florissi talked about the Netscript project [Florissi, 98] approach to programmable networks. Differentiating their approach from other Active Network projects by separating protocol code deployment from active execution, he introduced the notion of *virtual active networks (VANs)*. Danilo Florissi described a language-based Netscript network programmability approach with the analogy of postscript for printers. Layered over Java, the Netscript language and node/network environments serve as a foundation for building VANs on-the-fly. The speaker used a dataflow model for constructing packet stream processing engines over a network abstraction for end to end connectivity. Danilo presented an overview of the VAN architecture and the capability to coordinate composition of VAN units to deliver end to end bandwidth management, interoperability and protection. He then described a methodology for managing active networks using compiler generated MIBs and design-time manageability.

How Virtual Networks Can Support QoS-Sensitive Internet Applications, Jens-Peter Redlich, NEC, C&C Research Labs

The third speaker of the session, Jens-Peter Redlich, talked about support for QOS in Internet-based virtual networks [Redlich,98] through an *intelligent router*. The speaker discussed the multi-faceted problems facing users and service providers regarding performance, pricing and the evolution of Internet applications and IP networks. He proposed the intelligent router approach; describing the router architecture as serving virtual networks according to traffic policy. Jens-Peter described the intelligent router architecture as consisting of a policy database, billing mechanisms, a classifier which recognizes applications and user traffic. In addition, it included appropriate signaling mechanisms based on virtual network type (e.g., Integrated Services, Differentiated Services, native ATM and Active Network), and finally, a composition of schedulers, which perform resource allocation, admission control and predictive forwarding.

The XBONE: Building Overlay Networks, Anindo Banerjea, ISI

Anindo Banerjea delivered the final presentation of the session. Presenting the X-Bone Project [Banerjea,98], The speaker described the reasoning behind *overlay-based virtual network* as a simple means for containment and reservation. Some of the limitations of current overlay (e.g., M-Bone, 6-Bone) systems are that they are manual, administratively challenging and do not provide service guarantees. X-Bone seeks to remove the “lab coats” from the configuration complexity, characterized by Anindo, as offering a user-level, safe, rapid incremental deployment system over an IP-architecture. Our speaker mentioned that the X-Bone architecture, which is built on a foundation of IPv4 multicast, includes: i) an xd GUI, which announces and coordinates distributed overlay control over a two-level multicast channel; and ii) an overlay manager, which handles resource reservation through distributed resource daemons. Anindo Banerjea provided some examples of the X-Bone system and insights into the integration of virtual and physical network components. Finally, he concluded his talk by discussing static/dynamic overlays, application

routing, and his hope that X-Bone would serve to deploy, manage and integrate active networks.

Discussion

At the end of the session, the presenters formed an ad hoc panel to discuss the main themes of the session. The session chair asked what the panelists had in mind for QOS for virtual networks. Danilo Florissi said that the nodeOS [DARPA,96] should take care of it, while Anindo Banerjea suggested an orthogonal relationship. John Vicente from Intel said that well-grounded resource management techniques were essential to achieving QOS in virtual networks. In addition, John Vicente noted tunneled virtual networks alone would not fly in the long haul commercially. The “VN-QOS” would need to be solved. A related question posed by the audience asked whether implicit or explicitly signaling should be used in the context of future IP QOS. Paul Ferguson noted that both models would be required depending on the application and customer needs. The virtual network service within the context of the Internet was raised. John Vicente commented that there were already a number of virtual network solutions that automated tunneling on the market. Several panel members mentioned management, billing, video conferencing and IP telephony as obvious services for a virtual network capable Internet.

ACTIVE NETWORK INTERFACES

Chair: Ken Calvert, University of Kentucky

Active Reservation Protocol, Bob Braden, ISI

Bob Braden began his talk on Active Reservation Protocol (ARP) [Braden, 98] by making some general observations regarding the evolving nature of router software (e.g., highly complex and unwieldy updates) and static signaling protocols, e.g., RSVP. Protocol portability, dynamic versioning and user customizability are essential to overcoming the challenges facing the router software industry. Leveraging active network ideas, the speaker described a router/switch that included an active execution environment that enabled code portability, dynamic extensions and customization for protocol/version selection. Using alternative distribution mechanisms, incremental Java-based functional extensions employing class inheritance techniques can be dynamically extended from base versions. Bob

Braden discussed an active network project at ISI that uses an active execution environment with an ASP Exec module that implements the Protocol/Versions by intercepting signaling messages. The messages identify the VersionSpec class and load the appropriate VersionContext class according. The appropriate packet processing method is invoked followed by the loading of the functional extensions using the nodeOS.

Commercial Use of Active Networking, Dan Nessett, 3Com

The second speaker of the session, Dan Nessett, speculated [Nessett,98] about the commercial realization of active networking technology by discussing the commercial requirements of legacy integration, engineering support and business return on investment. The speaker made the argument that “shrink-wrapped” active technology with scalable extensions (i.e., not homegrown routers or stand alone services) as well as common, commercial-grade languages were essential. Dan Nessett presented an overview of an experimental execution environment and architecture that had been prototyped for investigation of an active capsule-based RSVP signaling system and active device management using LDAP and SNMP.

The Scout / Joust Active Network Node, Andy Bavier, Princeton University

Andy Bavier gave the final presentation [Bavier,98] on the Scout OS for network appliances (e.g., NetTV) and Joust, a configuration of Scout. He began by describing a key component of the Scout framework called the BERT scheduler. The scheduler based on a virtual clock packet switch algorithm supports both real-time and best-effort paths (a key Scout abstraction) for end to end scheduling and resource accounting. Andy Bavier described an IP router, TCP forwarder, and connection splicing abstractions, which leveraged a path abstraction and the Scout OS for end to end composition. Andy Bavier concluded by discussing the Joust JVM implementation using a number of example scenarios to illustrate active NetTV, active RSVP signaling and active proxy technologies.

Discussion

The panel discussed the role of active networks in future networks. Bob Braden thought its main role would be management and then network control in that order. Andy Bavier agreed with him and noting that forwarding should be done by hardware not Java. As all of the panelists mentioned Java in their presentations, but Bob Braden reminded us that Java has a license issue pending resolution.

PROGRAMMABLE NETWORKING

Chair: Gisli Hjalmtysson, AT&T Research

PRONTO: A Platform for Programmable Networking, Pawan Goyal, AT&T Research

In the first presentation of the Programmable Networking Session, Pawan Goyal argued for network programmability as a means for faster deployment of new services, more user control over resources and increased network flexibility. The speaker presented the PRONTO [Goyal,98] programmable platform that makes appropriate trade-offs between the control plane programmability and active packet forwarding. Through a kernel-level classifier, Pawan Goyal explained that the PRONTO controller manager enforces the appropriate *degrees of programmable control* through reservation protocols, frame peeking, queue manipulation and appropriate scheduling mechanisms. Pawan Goyal divided the PRONTO functional services. First, the kernel level services include signaling, classifier, frame peaking, CPU scheduler: H-SFQ, link scheduler components of the architecture. Next, user-level services include a controller manager and user installed program components. In addition, the speaker, discussed PRONTO applications, including selective discard, an active retransmission server, and flow-specific execution environments. The work represents a taxonomy based on the degree of required programmability that 'activate' the necessary PRONTO architectural components.

An Active Network Approach to Efficient Network Management, Yuval Shavitt, Lucent Bell Labs

Next, Yuval Shavitt spoke [Shavitt,98] on active network techniques for efficient network management. The proposed methodology involved distributing network management tasks, shortening control loops and reducing redundant management information. Through a

router diverter add-on, SNMP packets are diverted to an *active engine*, which generates sessions with the appropriate SNMP data consumer and manages (through an active manager) the data transfer and session usage profiles. The speaker argued that network management could be the ideal application domain for active networks and recognized security and safety as relevant factors in active network deployment.

Peter Steenkiste, An API for Runtime Network Resource Management, CMU

Next, Peter Steenkiste spoke about an API for runtime resource management [Steenkiste,98] based on the work at CMU on Darwin. He characterized resource management and QOS criteria as spanning across time, space an organizational dimensions. The speaker introduced the virtual network abstraction as an integrated set of resources with the requirement for customized resource management support. A conceptual framework of the Darwin architecture was presented describing resource brokers, delegates, a hierarchical scheduler and a signaling protocol. Peter Steenkiste explained that the Darwin resource management mechanisms were designed for fine (hierarchical scheduler), medium (delegates) and coarse (resource broker) network control and customization. Within the Darwin framework, Peter Steenkiste focused primarily on the dynamic network control through a delegate control API. Based on Java code segments and a flow network model, the delegate control mechanisms perform resource monitoring, local resource management, flow redirection and inter-delegate communications for end to end coordination and synchronization. He illustrated several examples where delegates were used to dynamically control customized selective frame dropping, adaptive encoding control, active flow management and load sensitive flow rerouting.

Composing Services in CANES, Ken Calvert, University of Kentucky

Ken Clavert closed the session with a presentation on composing services using CANEs [Calvert,98]. Active networking concepts and their enabling service benefits were briefly discussed. The granularity, virtual machine power and programmable authority are relevant aspects in the programmability equation. CANEs provides a methodology for

designing and implementing a “network API” for programmability. Building on this, he introduced the *slot programming model*, where user-defined behavior can be injected to processing slots with modularized programs and slot scheduled for execution; thus, allowing customization of network services. LIANE was then discussed. This model is used for service composition, program correctness and preservation in CANEs. Through simulation, experiments were performed with LIANE to observe application behavior for local congestion adaptation for multicast video, mobility and in-network caching techniques.

VIEWS OF DIFFERENTIATED SERVICES

Chair: Kathleen Nichols, Cisco

Diff Services: Magic Bullet or Cure for Cancer ?, John Wroclawski, MIT

John Wroclawski opened the diff-serv session with a presentation on “Diff Services: Magic Bullet or Cure for Cancer?” [Wroclawski,98]. Providing some background information on IP QOS, he talked briefly about int-serv directions and motivations discussing some issues that had plagued the adoption of the service model and RSVP. These “plagues” included topology complexity, scalability and administration, which he claimed were falsely overstated or without validation. The speaker went on to discuss the diff-serv goals clarifying some misconceptions on the effectiveness of the proposed approach. These included distinguishing signaling (e.g., RSVP) mechanism for resource allocation versus aggregate classification and forwarding mechanisms. John Wroclawski concluded with an observation: diff-serv scales well under typical topological scenarios, however, with increased network and domain complexity it may end up with similar “plagues” like int-serv.

Using Differentiated Services, Kathleen Nichols, Cisco

Next, Kathleen Nichols discussed the background behind the diff-serv initiative [Nichols,98]. The speaker described how the diff-serv aggregate behavior is marked and identified through packet codepoints and forward treatment using per-hop behaviors (PHBs) in the core network. Support for flow/microflow state is pushed to edge networks. Differential services are positioned as a building

block for adding new services by adding rules to behavior aggregates. There is a focus on standardizing the mechanism and not service. Services can extend through diff-serv network across domains with bilateral agreements. The speaker presented two examples, one enterprise network and the other an ISP network. An end to end “virtual leased line” concept was introduced based on an expedited forwarding PHB and appropriate policing and traffic conditioning to ‘emulate’ the notion of dedicated leased lines. A bandwidth broker architecture was also introduced as an appropriate allocation system for diff-serv routers, manage inter-domain bandwidth requests, and maintain domain policy and security.

Policy and Management Issues in a QOS-enabled Internet, Scott Bradner, Harvard University

Scott Bradner spoke on policy and management in a QOS-enabled Internet [Bradner,98]. Policy in QOS was necessary to insure appropriate allocation of resources and service quality billing. The solution must address appropriate network engineering including reservation, class policing and traffic optimization. In addition, any solution should also encompass network monitoring including class traffic measurement and reservation service quality. Scott Bradner also spoke about policy based on “AAA” for usage *authority*, requestor *authentication* and *accounting* for billing. He argued that class-based and flow-based service delivery measurements are essential for customer paid services. The implementation would require AAA servers and per-domain bandwidth brokers as described by Kathleen Nichols.

Two-Tier Resource Management, Fran Reichmeyer, Nortel Networks

Fran Reichmeyer presented the final talk [Reichmeyer,98] of the session on a two-tier model for Internet QOS resource management. QOS resource management strategy should parallel the hierarchical routing architecture existing in the Internet to address scaling and administrative control issues properly. Fran Reichmeyer speculated that end to end QOS could be achieved through intra-domain policy, inter-domain SLAs and concatenation of bilateral agreements. The speaker emphasized the use of the bandwidth broker for inter-domain communication (e.g., SLA negotiation),

aggregate provisioning, intra-domain resource management and diff-serv treatment at domain borders. Fran Reichmeyer proposed to combine (two-tier) RSVP signaling with the diff-serv model for delivering end to end reservation and QOS/SLA delivery across domains requiring quantitative QOS. He argued for standard interoperability signaling protocols, which can also be used to signal resource management through policy services, e.g., COPS.

Discussion

Several people raised the question that the QOS problem was fundamental and IP or diff-serv IP did not necessarily have magic solutions. Differential service covers a wide range of technology and provision for “simple-QOS” may be within the diff-serv charter but the difficult problems of how the reservation and the bandwidth brokers work to support QOS is unclear. In addition, more complex services, (e.g., end to end QOS, on-demand bandwidth guarantees and billing) will require mechanisms above and beyond what has already been discussed. A member of the audience suggested that diff-serv may not stay simple in the end. Perhaps programmability could ease its evolution?

MOBILE AGENTS & OPEN SIGNALING

Chair: Roger Imprey, National Research Council

Tunnel Agents for Enhanced Internet QOS Orazio Tomarchio, Universita’ di Catania

Orazio Tomarchio began the session with a talk on tunnel agents [Tomarchio,98]. Orazio Tomarchio proposed a mobile agent paradigm as an alternative to QOS provisioning. A wide range of applications including distributed computation and management, mobile computing and application services require on-going interaction rather than ongoing communications. Orazio Tomarchio introduced a mobile agent platform for developing and managing mobile agents. The speaker focused on two key applications of the work, network management and QoS management. User, system, network and service agents can be implemented to achieve intelligent and distributed configuration management, fault management and service provisioning. Orazio Tomarchio characterized inadequacies in current partial reservation (with non-RSVP capable)

configurations, and described how *tunnel agents* can be setup at the edges of the network to enhance the reservation scheme.

Managing Personal Mobility using Software Agents, Ramiro Liscano, National Research Council

Next, Ramiro Liscano spoke [Liscano,98] about managing personal mobility. Focusing on application-level services, he presented a framework for personal communication agents to support “any-type” devices for manipulating user services and personalized information. The speaker presented an agent model, which consisted of an agent UI, JESS (knowledge base & reasoning), KQML (agent and user communications) and Java and CORBA execution environments. The speaker discussed heterogeneous user communication models and personal agent mobility management systems.

Joint Submission to OMG Wireless Jit Biswas, Kent Ridge Digital Labs

Jit Biswas presented [Biswas,98] a response to the OMG Telecom DTF Wireless RFI. The speaker described a wireless ATM environment with computational and bandwidth problems inherent in wireless environments. In addition, he described some design choices (e.g., distributed programming methodologies) and issues (e.g., channel characteristics, loss, delay, memory footprint) that must be faced in wireless signaling of mobile terminals. Selecting the CORBA “bus” model for object services (e.g., real-time message delivery, event and notification services, naming, trading, location) he stressed the importance of these CORBA interfaces in the area of wireless transport and signaling.

IEEE P1520 Standards Project: Application Programming Interfaces for Networks, Jit Biswas, Kent Ridge Digital Labs

Jit Biswas switched hats for the next talk and emphasized the need for standards activities, including the IEEE P1520 project for network programmability. Jit Biswas presented an overview of the on-going work in the IEEE PIN initiative and discussed the scope of its three working groups. The next PIN meeting will be held at Columbia University at the same time as INFOCOM and OPENARCH’99; see [P1520, 98b] for details.

How to Protect Own Interests in Open Communication Environment
Serguei Mankovski, Mitel Corp.

Serguei Mankowski gave the final presentation [Mankowski,98] of the session, he contended that protection of individual interests may be violated in communications over open environments through feature interactions. Well known in the telephony environment, it has been dealt with through exposure of low-level switching details to the application services. The speaker argued that network security is a feature interaction problem and discussed issues for open signaling infrastructure that can help protect the individual party interests.

OPENSLOT SESSION

Chair: Nikos Anerousis, AT&T Labs

The OPENSLOT Session provided fifteen minutes to participants who wanted to voice an opinion. First up was Raju Rajan (IBM) who provided some insights on network policy rules and administration [Rajan,98]. Raju Rajan discussed how policy rules can be used to achieve security and QOS. Next, Livio Ricciulli (SRI) spoke about the Anetd active network daemon [Ricciulli,98]. The speaker expanded on the Anetd model through service deployment and security mechanisms and described a set of commands used to manage an experimental overlay active network. Following this, Sanjai Narain (Bellcore) [Narain,98] talked about the need for "making network services easier to use" through service intelligence at the edges of the network. The speaker described some of the problems associated with configuring and troubleshooting Internet access service demonstrating a proactive utility called "Dr. Dailup" that resolves end to end component configuration, state and connectivity. Our next speaker was Dan Ionescu (University of Ottawa) who spoke on a comprehensive framework for network management [Ionescu,98] using CORBA. The final presenter of OPENSIG'98 was Chris Edwards (Lancaster University) who talked about an *open service model (OSM)* [Edwards,98] that provides mechanisms for application programmers to specify control requirements in different contexts and granularity. Through a lightweight meta-signaling protocol and layered, open interfacing/binding abstractions, the application

developer can enable multiple control mechanisms to manage end to end requirements.

CLOSING REMARKS

Irene Katzela (Chair) closed the workshop with a few final remarks. Irene asked, "after two days of discussion are we less confused than before?" and answered, "no!". A number of new themes emerged at OPENSIG'98 raising many new questions about where the whole programmable network movement was heading. If the number of questions was a reflection of new challenges ahead then we have a promising future. Irene reflected that she thought we have just scraped the top of the iceberg.

NEW MOVES IN 1998

It has been a strong year for new developments in the area of open programmable networking. A few of the "new moves" this year include:

The IETF established (November, 98) a *GSMP* [GSMP,96] Working Group that will focus on developing a general switch management and control protocol for IP routers. Peter Newman (GSMP architect) is an OPENSIG Alumni who presented his proposal on GSMP for ATM switches at an early OPENSIG meeting.

Another interesting move came with the formation of the *Multiservice Switching Forum (MSF)* [MSF,98] (formed in November 1998). The MSForum is an industry forum committed to an open, standardized control interface that will support multiple control planes for ATM systems. MSF promotes the separation of control and user/data plane for ATM-capable switches and open intra-switch interfaces across switch network provider platforms for heterogeneous service provisioning. The MSF is discussing the switchlet [Cplane,98] concept and GSMP. The switchlet came from Cambridge University who play a key role in OPENSIG.

Finally, the *Parlay Group* [Parlay,98] recently introduced an API specification which gives network customers the ability to create their own dynamic telecommunications applications.

OPENSIG INFO

All past OPENSIG proceedings are maintained electronically and are available at [OPENSIG, 96]. An "opensig-announce" e-mail exploder has been set up for sending out general

announcements pertaining to network programmability. The address is opensig-announce@comet.columbia.edu. To subscribe to this mailing list, please send e-mail to opensig-request@comet.columbia.edu.

OPENSIG'99

Co-Chairs for OPENSIG'99 are Peter Steenkiste and Hui Zhang (Carnegie Mellon University). For details on the next workshop see:

www.cs.cmu.edu/~cmcl/opensig99

ORGANIZED BY

Communications Group,
University of Toronto

PROGRAM COMMITTEE

Chair: Irene Katzela, University of Toronto

Anindo Banerjea, ISI
Andrew T. Campbell, Columbia University
Simon Crosby, and University of Cambridge
Aurel A. Lazar, Columbia University
Ian Leslie, University of Cambridge

LOCAL ORGANIZATION COMMITTEE

Kitty Chan, University of Toronto
Raymond Liao, Columbia University

REFERENCES

[ANmeets] DARPA Project Meetings,
www.darpa.mil/ito/research/anets/meetings.html

[Banerjea,98], Anindo Banerjea (USC/ISI), "The XBONE: Building Overlay Networks", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Bavier,98] Andy Bavier, (Princeton University), "The Scout / Joust Active Network Node", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

ACKNOWLEDGEMENTS

We would like to thank the Program Committee for their efforts in making this a successful workshop. Our special thanks to the Communications Group at the University of Toronto for hosting the event and to the Local Organization Committee of Kitty Chan (University of Toronto) and Raymond Liao (Columbia University). Finally any errors or omissions are ours and for these, our apologies in advance.

[Biswas,98a] Jit Biswas, (Kent Ridge Digital Labs), "Response to Request For Information – Telecom Domain Task Force, OMG" *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Biswas,98b] Jit Biswas, (Kent Ridge Digital Labs), "IEEE P1520 Application Programming Interfaces for Networks: An Introduction", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Braden,98] Bob Braden, (USC/ISI), "Active Reservation Protocol", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Bradner,98] Scott Bradner, (Harvard University), "Policy and Management Issues in a QoS-enabled Internet", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Calvert,98] Ken Calvert, (University of Kentucky), "Composing Services in CANEs", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Cplane,98] Control Plane Technologies Inc
www.cplane.com/

[Diff,98] Differential Services Working Group,
www.ietf.org

[DARPA,96] DARPA Active Nets Program,
www.darpa.mil/ito/research/anets/projects.html

[Edwards,98] Chris Edwards, (Lancaster University), "Meta-Signalling for Open Network Control", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Ferguson,98] Paul Ferguson (Cisco), "What is a VPN", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Florissi,98] Danilo Florissi (Columbia University), "The NetScript Approach to Programmable Networks", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Genesis,98] Programmable Virtual Networking, comet.columbia.edu/genesis/

[Goyal,98] Pawan Goyal, (AT&T Research), "PRONTO: A Platform for Programmable Networking", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[GSMP,96] General Switch Management Protocol from Peter Newman (Netsim Inc www.ensim.com/) find the specification at <http://www.iprg.nokia.com/protocols/rfc1987.htm>

[Ionescu,98] Dan Ionescu, (University of Ottawa), "A Corba Environment for Network Management: Results and Perspectives", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Liscano,98] Ramiro Liscano, (National Research Council), "Managing Personal Mobility using Software Agents", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Mankowskii,98] Serguei Mankowskii, (Mitel Corp), "How to Protect Own Interests in Open Communication Environment", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Mihai,98] Nelu Mihai (AT&T West Labs), "Geoplex: An Open Service Platform", Keynote Address. *Proc. 5th Workshop on Open Signaling*

for ATM, Internet and Mobile Networks, Toronto, October 5-6, 1998.

[Mobiware,98] Full Source Code Distribution for Open Programmable Mobile Networking comet.columbia.edu/mobiware

[MSF,98] Multiservice Switch Forum www.msforum.org/

[Narain,98] Sanjai Narain, (Bellcore), "Making Network Services Easier to Use", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Nessett,98] Dan Nessett, (3Com), "Commercial Use of Active Networking", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Nichols,98] Kathleen Nichols, (Cisco), "Using Differentiated Services", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[OPENARCH,97] International Conference on Open Architectures and Network Programming, comet.columbia.edu/openarch

[OPENSIG,95] First Working Group Meeting, October 24, 1995, Columbia University, <http://comet.columbia.edu/opensig/background.html>

[OPENSIG,96] OPENSIG Working Group <http://comet.columbia.edu/opensig/>

[P1520,98a] IEEE Standard for Application Programming Interfaces for Networks, www.ieee-pin.org/

[P1520,98b] Biswas,J., et al., "Application Programming Interfaces for Networks", IEEE P1520 Working Group Draft White Paper, <http://www.ieee-pin.org/>

[Parlay,98] <http://www.parlay.org>.

[Rajan,98] Raju Rajan, (IBM T.J.Watson Research Center), "Policy Rules! Network Policy Administration for Security and QoS", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Redlich,98], Jens-Peter Redlich (NEC), "How Virtual Networks can support QoS-sensitive

Internet Applications", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Reichmeyer,98] Fran Reichmeiyer, (Nortel Networks), "Two-Tier Resource Management", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Ricciulli,98] Livio Ricciulli, (SRI), "Anetd: Active Network Daemon", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Shavitt,98] Yuval Shavitt, (Lucent Bell Labs), "An Active Network Approach to Efficient Network Management", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Steenkiste,98] Peter Steenkiste, (Carnegie Mellon University), "An API for Runtime Network Resource Management", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Tempest,98] Cambridge University Tempest Project www.cl.cam.ac.uk/Research/SRG/dcan/

[Tomarchio,98] Orazio Tomarchio, (Universita' di Catania), "Tunnel Agents for Enhanced Internet QoS", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Vicente,98] John Vicente (Columbia University & Intel Corporation), "The Genesis Project: Toward Programmable Virtual Networking", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998. <http://comet.columbia.edu/genesis/>

[Wroclawski,98] John Wroclawski, (MIT), "Diff Services: Magic Bullet or Cure for Cancer?", *Proc. 5th Workshop on Open Signaling for ATM, Internet and Mobile Networks*, Toronto, October 5-6, 1998.

[Xbind,96] Code Distribution for Project Xbind, <http://comet.columbia.edu/xbind/>

[Xind,98] Xind Inc. www.xbind.com/