Report on DIMACS Workshop on Quality of Service Issues in the Internet

February 8-9, 2001

Workshop Organizers:

Funda Ergun, Case Western Reserve University
Michael Mitzenmacher, Harvard University
Cenk Sahinalp, Case Western Reserve University
Aravind Srinivasan, Bell Labs
Bulent Yener, Bell Labs

1 Workshop Focus

This workshop was aimed at identifying the fundamental problems associated with providing QoS in the IP networks, both from the theoretical and application-oriented points of view The explosive growth of the Internet is accompanied by two paradigms. First, the IP-based networking is expected to be the dominant infrastructure. Second, unconventional applications such as IP telephony, Virtual Private Networking, Real-time transactions, IP HDTV, are expected to be carried on the IP-based networks. These two paradigms are translated to the following challenge: how to provide Quality of Service (QoS) in the IP-based Internet. Fundamentally the IP-based networking is designed for delivering data traffic with best-effort (i.e., no guarantee) service; thus it is not capable of providing end-to-end QoS. The challenge is to devise new protocols and algorithms for providing QoS while maintaining the basic principles of IP-based networking.

Topics covered in the workshop included (1) QoS on IP networks, both wired and wireless; (2) emerging high-speed technologies of QoS; (3) IP over ATM, IP over SONET, etc.; (4) What is QoS? (Requirements, resources, allocation/protection of the resources; wireless IP networks; QoS routing; packet filtering/classification; QoS resource reservations/signaling; congestion control; multilayered multicast; and QoS for network storage/caching).

2 Summary of the Presentations

The following are author-supplied abstracts, mildly edited.

Effects of Filler Traffic in IP Networks
Nasser Alzidi, Adam Feldman, Vincenzo Liberatore Case Western Reserve University

On the Internet, there is a well-documented requirement that much more bandwidth be available than is used on average. This extra bandwidth is required to handle "network spikes", those times when network traffic peaks, using much more bandwidth than is normally needed. Sometimes these spikes can be quite large, requiring a network be provisioned to provide two to twenty times the bandwidth than it uses under normal circumstances in order for the network to maintain its adequate speed of operation and prevent degradation for its users. Thus, for a network to achieve acceptable performance a good portion of its bandwidth must sit idle most of the time. Due to the oft high prices of bandwidth, this is either not feasible or leads to quite a waste of money. However, if this extra bandwidth could be used for other activities while it is not needed to handle a spike of traffic, then the network would be much more cost effective. The goal of this project is to determine a method of allowing the extra bandwidth to fill its idle time by carrying low priority traffic. When the network spike occurs, normal traffic would take priority over this "filler" traffic, allowing the network to use this bandwidth as if it had been idle. Since the filler traffic is composed of non-time sensitive data, the interruption the spike causes is expected and not disruptive. This way the network is equipped to handle all spikes of traffic without slowdown, while at the same time, the excess bandwidth is being used to accomplish work.

The problem in executing this idea is deciding the nature and specifics of the filler traffic in such a way that prevents it from interfering in any way with the normal, or pre-existing, traffic, yet ensures that it will eventually reach its destination. To see the effects different types of traffic have on the pre-existing traffic, the results of a network simulation, using traffic data from a Harvar trace and completed with NS-2, were studied. The simulation was executed many times for the same input data, to see the effects of changing different parameters, such as size and type of filler traffic (CBR or FTP), filler buffer, and bandwidth and latency of the central link of the simulated network. The results of this experiment are packet dynamic, breaking down into four items, which were charted over the experiments, for a given changing parameter: average packet delay, percent of dropped packets, ack compression, and the amount of bandwidth used. Each of these result categories were plotted separately based on which type of traffic (pre-existing and filler) and which direction, to show how the filler traffic in one direction affects the pre-existing traffic in both directions.

The resulting output charts indicate that the idea of implementing filler traffic is feasible. They also confirm that the exact parameters used can have a profound impact on the effect the filler traffic has on the pre-existing traffic. If the filler packets are too large or sent at too high a rate, then the pre-existing traffic is slowed. On the other hand, the usefulness of the filler traffic is also dependent on the exact parameters used. For example, if the filler buffer is very small, then the majority of the filler packets will be dropped, resulting in little or no work being done by having the filler traffic.

The next step is to study the effects of this filler traffic across a network with a low bandwidth delay product (BDP), such as that of a bank of modems. Since a modem network is much slower than a high-speed local network, there is potential for filler traffic to adversely affect the former even though the latter is unaffected. Once it is determined how these two very different types of networks are affected by each type of filler traffic, a general scheme of using specific filler traffic can be developed and tested for effectiveness.

Fast and efficient bandwidth reservation for robust routing

Lisa Fleischer, Carnegie-Mellon University

Given a set of point-to-point bandwidth demands, and a network with fixed bandwidth capacity, we examine the problem of reserving sufficient bandwidth in the network to route all demands, even in the case of node or link failures. In addition, we wish to enforce that traffic is sent over paths with small number of hops. Thus we try to minimize the sum of bandwidth reserved over all links.

One approach is to insist on dedicated reserve bandwidth for each commodity. This does not allow full utilization of network bandwidth. Our approach is to allow reserve bandwidth to be shared by multiple commodities. Experimental evidence indicates that this allows for much higher network traffic volumes.

Our problem may be formulated as a linear (or integer) program. However, this linear program is very large, and thus for real routing scenarios, it may be impractical to devote sufficient time to solve the LP exactly.

We develop an FPAS for solving the linear program. While this problem is a multicommodity problem, it does not fall into the framework of existing FPAS's for multicommodity flow due to the fact that we need to reserve extra bandwidth on separate paths in case of failures. Thus we extend existing techniques to solve our problem. To obtain integer solutions, we round the fractional solution from the FPAS and then use local improvement.

Co-authors: A. Meyerson, I. Saniee, A. Srinivasan

Delay-sensitive unicast and multicast routing
Ashish Goel, University of Southern California

With the growing popularity of delay sensitive services such as voice and video, it is becoming more important to have efficient routing schemes that result in routes that have a low delay as well as a low-cost. In this talk, we will address this problem for both unicast and multicast routing.

For unicast routing, we will present a series of engineering observations that lead to an efficient algorithm for computing delay constrained shortest paths from one source to all destinations. For multicasting routing, we will present an online algorithm that is $O(logn)$ competitive with respect to the total cost of the multicast tree, while guaranteeing, at the same time, that the distance from the source to any receiver along the tree is at most a constant factor more than the shortest path distance. We will also describe our preliminary implementation/simulation results for these algorithms.

The unicast routing algorithm represents joint work with Ramakrishnan, Kataria, and Logothetis. The multicast routing algorithm represents joint work with Munagala.

On The Impact of Aggregation on The Performance of Traffic Aware Routing
Roch Guerin, University of Pennsylvania

This paper investigates the impact of traffic aggregation on the performance of routing algorithms that incorporate traffic information. The focus is on two main issues. Firstly, we explore the relationship between the *average* performance of the network and the level of granularity at which traffic can be assigned to routes. More specifically, we are interested in how average network performance improves as the ability of the routing protocol to split traffic arbitrarily across multiple paths in the network increases. Secondly, we focus on the impact of traffic aggregation on short-term routing behavior. In particular, we explore the effects of traffic aggregation on traffic variability, which directly affects short-term routing performance. Our analysis is based on traffic traces collected from an operational network. The results of this study provide insights into the cost-performance trade-offs associated with deploying routing protocols that incorporate traffic awareness.

This is joint work with A. Sridharan, S. Bhattacharyya, C. Diot, J. Jetcheva, and N. Taft

Scalable Resource Reservation for Internet QoS
Ellen L. Hahne, AT&T Labs

A growing number of Internet applications could benefit by reserving network resources like link bandwidth. Such applications include real-time voice and video, better-than-best-effort data services, virtual private networks, and traffic engineering.

The current Internet QoS architecture, called IntServ, with its signalling protocol RSVP, operates per-flow and end-to-end. Consequently, it suffers from three problems: high packet forwarding costs, excessive protocol overhead, and administrative difficulties when operated across domain (i.e., carrier) boundaries. Since the number of packets, flows, and domains in the Internet are increasing rapidly, these scaling issues must be addressed soon if Internet QoS is to become a reality.

A newer proposed Internet QoS architecture, called DiffServ, solves the packet forwarding scalability problem, by classifying packets into a tiny number of classes at the edge of the network. But the problems of inter-domain administration and reservation protocol overhead remain. For INTRA-domain resource reservation, where protocol scalability is less of a concern, RSVP may be appropriate. Other intra-domain solutions are also possible; e.g., generously engineered networks may not require resource reservation at all. Intra-domain choices can be left to the individual carrier. But for INTER-domain resource reservation, we propose a new protocol, called the Border Gateway Reservation Protocol (BGRP). BGRP is highly aggregated and scales well, in terms of message processing load, state storage and bandwidth. We believe that this combination – DiffServ in the data plane, carrier's choice (e.g., RSVP) for intra-domain control, and BGRP for inter-domain control – is the right approach for Internet resource reservation.

BGRP works as follows. It is designed for unicast traffic, and it builds distributed data structures in the form of sink trees. There is one sink tree for each destination domain and each DiffServ traffic class. A sink tree aggregates reservations of a given traffic class from all data sources in the network to the given destination domain. Since backbone routers maintain only the sink tree information, the total number of reservations at each router scales linearly with the number of Internet domains N. (Even aggregated versions of RSVP have a reservation count that can grow like $O(N^2)$.) We will argue that BGRP introduces precisely the right degree of aggregation for the current and future Internet.

This is joint work with Ping Pan (Juniper Networks) and Henning Schulzrinne (Columbia University).

Dynamic Routing of Restorable Bandwidth Guaranteed Paths in MPLS Networks
T.V. Lakshman, Murali Kodialam, Lucent Technologies, Bell Laboratories

This talk will focus on algorithms for the routing of bandwidth guaranteed paths in MPLS networks that survive single link/node failures. Unlike traditional network design problems, where all the requests are known ahead of time, in this talk we assume that bandwidth requests come one at a time and have to be routed at the time of arrival. Each request is routed along an active path with a designated backup path that will be used to restore the connection in the case of network element failures. Several different restoration options will be explored. The relative advantages of the different options and the algorithmic implications will be outlined.

Effects of Pricing on Nash Equilibria and Total User Utility
Andrew Lomonosov, Meera Sitharam, Department of Computer and Information Science and Engineering University of Florida, Kihong Park, Department of Computer Sciences, Purdue University

This talk deals with Quality of Service provision problem in noncooperative networks where applications of users are selfish. We study a model of QoS provision in noncooperative networks where every user is required to use at most one service class. Individual user utility determines level of satisfaction that user receives per unit of his traffic volume from a certain Quality of Service and a price charged for this unit of volume. Utility is inversely proportional to price and directly proportional to Quality of Service. Both Quality of Service and pricing per unit $p()$ are decreasing functions of a total volume q in the service class. Thus individual user utility is a function $U : q \rightarrow R$ from class volumes to the nonnegative reals. We assume that when QoS is less than certain threshold user derives no utility, hence

$U(q) = 0$ if $q > b$. When $0 \leq q \leq b$ utility is a nondecreasing function.

We show that if $p()$ is a constant function then a Nash Equilibrium will always be reached starting from any initial configuration, using any sequence of selfish moves by individual users.

If $p()$ is a non-constant increasing function then we show a similar result if all selfish moves improve the utility to a moving user, while not exceeding thresholds of other users.

If $p()$ is a non-constant increasing function and selfish moves can exceed thresholds of other users then Nash Equilibrium is not always reached. We examine the effect various non-constant pricing functions have on total user utility as well as on the existence of Nash Equilibrium. We show that a random increasing pricing function is likely to improve total user utility, however it is not likely to induce a Nash Equilibrium. We also show that certain special class of pricing functions always induces a Nash Equilibrium and is likely to improve total user utility.

Internet Traffic Managers
Ibrahim Mattam, Boston University

The desire to tame the unpredictability of today's Internet has led to recent efforts to embed rich functionalities into networking devices (e.g., switches and routers) in order to support specific applications (e.g., streaming media) and emerging paradigms (e.g., multicast). In tandem with these current trends, we believe that "Traffic Managers" (TMs)—special network elements strategically placed in the Internet—should implement control strategies for reducing traffic burstiness and improving network utilization.

This talk reviews some basic control capabilities that these TMs should employ to ensure desirable properties (e.g., satisfaction of delay requirements, compliance with TCP semantics, or improved fairness across flows). These control capabilities include differentiated control of flows with divergent characteristics, and aggregate control of "congestion-equivalent" flows. The focus of this talk is on the control of TCP flows as they constitute the majority of the traffic volume on the Internet today.

Based on control-theoretic principles, we illustrate the utility of a low-cost differentiated control applied through a TM, where TCP flows are "isolated" into classes based on their lifetime/size. This isolation provides a firewall between flows with divergent characteristics, including burstiness of arrivals/departures and congestion/sending window dynamics.

Integrated QoS Support in the Dynamic Mobility Agent (DMA) Architecture
Archan Misra, Telcordia Technologies

Our presentation discusses our architecture for integrating QoS assurances and mobility management in future IP-based cellular networks. We describe a Differentiated Services-based QoS architecture for next-generation wireless networks. The architecture is based on the two-level Dynamic Mobility Agent (DMA) hierarchical mobility management scheme and integrates Bandwidth Broker-based admission control and resource provisioning for mobile nodes. The DMA architecture is extended to satisfy the QoS requirements of a mobile node, while requiring it to specify its traffic profile only when it first moves into a domain. We explore alternative approaches for dynamically assigning Mobility Agents to a mobile node and evaluates their suitability for different service differentiation models. We also explain the advantages of defining a hierarchical diffserv profile enforcement mechanism in the cellular domain, especially from the standpoints of scalability and signaling efficiency.

Co-Authors: Sajal Das, Subir Das and Anthony Mcauley

Traffic Engineering and Control for the Management of Service Level Agreements in Packet Networks
Debasis Mitra, Lucent Technologies, Bell Laboratories

The talk considers templates for QoS-centered Service Level Agreements (SLA), and a framework for their real time management in multiservice packet networks. The SLAs are for QoS-assured delivery of aggregate bandwidth from ingress to egress nodes, where the elementary entities are flows or calls of various QoS classes. A SLA monitoring scheme is presented in which revenue is generated by the admission of flows into the network, and penalty incurred when flows are lost in periods when the service provider is not SLA-compliant. In the SLA management scheme proposed here the results of a prior design are used, in conjunction with measurements taken locally at ingress nodes, to classify the loading status of routes as either undersubscribed or oversubscribed. Our resource management is based on Virtual Partitioning and its supporting mechanism of Bandwidth Protection, which aims to carry as much of the offered flows at any time as is compatible with the protection of future undersubscribed routes. The effectiveness of SLA management is measured by the robustness in performance in the presence of great diversity in actual traffic conditions. The talk will describe a simulation testbed called D'ARTAGNAN from which numerical results for a case study are presented. The results show that the SLA management scheme is robust, fair and efficient over a broad range of traffic conditions.

Comfortable Service Based on Policed Priority Queuing for Per-Flow QoS Guarantee
Yasuo Okabe, Yasuo OKABE Graduate School of Informatics, Kyoto University

In order to guarantee QoS of each flow, Policed Priority Queuing (PPQ), which processes packets in constant time, and statistically guarantee bandwidth, packet loss ratio and queuing delay, is proposed. The behavior of PPQ is analyzed by elementary queuing theory. Comfortable Service (CS), which is based on PPQ and realizes statistical end-to-end QoS guarantee on the Internet, is discussed.

PPQ is already implemented in a hardware router, which can handle 1000 individual flows at a gigabit-class interface.

In an experimental environment using the routers, 15 video servers transmit 180 motion JPEG streams, each of which is 320x240 size, 30 frame per second, and 4.5 Mbps, on an one-Gbps interface. The total bandwidth of the streams is 810 Mbps. A client can replay a full-motion video stream without being influenced not only best-effort traffic, but also streams which violate the pre-assigned bandwidth of 4.5 Mbps, under the condition that the replayed video stream obey the pre-assigned bandwidth.

Voice QoS Metrics for Packet Based Communications Systems
Dan Quinlan, Lucent Technologies, Bell Laboratories

During the development of the wired analog network, engineering teams recognized from the start that voice quality was a primary design driver. As such, projects aimed at understanding speech mechanisms and the human auditory system were a key component of the telephony research executed in the first half of this century. The current migration from TDM to new packet-based (and wired to wireless) voice communication systems has given rise to a new set of signal degradations, and these impairments can have severe perceptual effects. While the ideal system would transmit voice at high bit-rates and with no signal degradation, packet voice systems are often designed using signal compression in the edge equipment, and the transmission paths are far from ideal. In this arena, mainstream communications systems can no longer be thought of as linear and time-invariant. So, many of the techniques used to evaluate performance are no longer workable, and now the communications industry is struggling to figure out how to systematically evaluate the effects of these new impairments. The first step is a recognition that a connection must be made between objective specifications and subjective response. The talk will present a broad overview of the key issues, and will present some thoughts regarding how the communications industry is responding to these problems.

Joint work with J. McGowan, T. Spencer, S. Pennock, Lucent Technologies, Bell Laboratories

A Simple Efficient Approximation Scheme for the Restricted Shortest Path Problem with Applications to QoS Routing in the Internet
Danny Raz, Bell Laboratories and Technion

In this talk I'll describe a very simple fully polynomial approximation scheme for the restricted shortest path problem. The complexity of this ϵ-approximation scheme is $O(|E|n(\log\log n + 1/\epsilon))$, which improves Hassin's original result [Has92] by a factor of n and is compared favorably with the much more complex result of Cynthia Phillips [Phil93]. Furthermore, this complexity bound is valid for general graphs regardless of the cost values.

I'll then describe how to apply these techniques to Quality of Service (QoS) routing in the Internet. It turns out that the problem of how to provide the required QoS level with a minimal cost is a natural generalization of the restricted shortest path problem. For example, one can use similar approximation schemes to efficiently compute the path that ensures a certain delay bound (say 200 ms) for voice over IP, with almost a minimal cost to the customer.

This is joint work with Dean Lorenz, Ariel Orda, and Yuval Shavitt

Minimizing Expected Service Time for Databroadcast
Nicolas Schabanel DIMACS and AT&T

Databroadcast protocols have been designed to unload the networks from delivering the most popular information. Instead of sending these informations thousands of times as soon as a user requests them, a databroadcast server continuously broadcasts these popular informations on a shared medium (such as wireless, satelite, ethernet,...): a user who is interested in one of these informations, connects to the media, monitors the channels until the information is broadcast. The server does not know the effective user requests but knows the popularities $p_1, ..., p_m$ of each information $M_1, ..., M_m$. The goal is then to find a schedule that optimizes the quality of service, that is to say the expected service time for a random user.

Several variants of this problem exist. We will present the latest results in this area including the study of the cases where the information have equal and not equal lengths and where preemption is allowed. This work is essentially algorithmic.

At the end of the talk we will also present possible extensions: What happens if we try to take into account the effective requests? What happens if one introduce dependencies between the messages?...

Buffer Overflow Management in QoS Switches
Baruch Schieber, IBM T.J. Watson Research Center

We consider two types of buffering policies that are used in network switches supporting QoS (Quality of Service). In the FIFO type, packets must be released in the order they arrive; the difficulty in this case is the limited buffer space. In the "bounded-delay" type, each packet has a maximum delay time by which it must be released, or otherwise it is lost. We study the cases where the incoming streams overload the buffers, resulting in packet loss. In our model, each packet has an intrinsic value; the goal is to maximize the total value of packets transmitted.

Our main contribution is a thorough investigation of the natural greedy algorithms in various models. For the FIFO model we prove tight bounds on the competitive ratio of the greedy algorithm that discards the packets with the lowest value. We also prove that the greedy algorithm that drops the earliest packets among all low-value packets is the best greedy algorithm. This algorithm can be as much as 1.5 times better than the standard tail-drop policy, that drops the latest packets.

In the bounded delay model we show that the competitive ratio of any online algorithm for a uniform bounded delay buffer is bounded away from 1, independent of the delay size. We analyze the greedy algorithm in the general case and in three special cases: delay bound 2; link bandwidth 1; and only two possible packet values.

Finally, we consider the off-line scenario. We give efficient optimal algorithms and study the relation between the bounded-delay and FIFO models in this case.

This is joint work with Alexander Kesselman, Zvi Lotker, Yishay Mansour, Boaz Patt-Shamir and Maxim Sviridenko.

A Distributed QoS Multicast Routing Protocol
Yuval Shavitt, Bell Labs and Tel Aviv University

Multicast is an efficient way to deliver content to a large group of receivers by using a tree structure along which all communication is carried out. Building a multicast tree, termed multicast routing, was studied extensively. Finding an optimal multicast tree, i.e., a tree that uses a minimal number of links, is known to be NP-hard. Thus heuristics were deployed to find trees that are "good enough". The simplest heuristic was to use the existing unicast routing paths and build a tree of shortest paths from the sources (or some core node in case of several sources) to the destinations. This approach was adopted by the IETF in its multicast standardization effort, where shortest path trees are used as multicast trees: DVMRP[Dee89], PIM [EFH+98], MOSPF[Moy94].

The shortest path trees work well for delivering information that is not QoS sensitive. However, when applications have non-trivial QoS Requirements (delay bound, throughput, etc.) there is a high probability that the shortest path trees would not have adequate characteristics [CNS00].

Let a feasible tree be a multicast tree that satisfies the QoS requirement. A feasible tree branch (path) is a path that connects a new group member to a multicast tree and has the resources to support the required QoS. The task of QoS multicast routing is to find feasible tree branches for new group members. Finding feasible tree branches is difficult in very large networks such as the Internet because it is impractical to maintain the global QoS state at any single node. A brute-force flooding algorithm that searches all possible paths in the network guarantees to find a feasible branch if one exists. However, the excessive overhead of full-scale flooding deems to be impractical for all but small networks. Thus, for applications that require QoS guarantees, recent research focuses on distributed multicast routing algorithms that search a selected subset of the network to find feasible tree branches for new group members [Sha96, CC97, CRS97, FPB98, CN99, CNS00].

The spanning join protocol by Carlberg and Crowcroft [CC97] relies on limited distance flooding to find a feasible tree branch to connect a new member. It starts with a small flood radius and use increasing larger radii after a failure to find a feasible path. While it works relatively well for small networks, its excessive communication overhead deems it impractical for large networks. QoSmic[FPB98] introduces an extra control element, called Manager router, to handle the join requests of new members. In large multicast trees the search for 'good' in-tree nodes to start the search of a path to the new leaf is fairly expensive.

A good QoS routing protocol achieves a favorable tradeoff between the routing overhead and the ability of finding a feasible path (often quantified as success ratio or success probability). In addition, a good protocol exhibits or optimizes other characteristics, such as minimizing the extra state information the protocol maintains in the network; decentralizing the routing operations to spread the workload; and adapting the routing activity according to the current network condition and avoiding the area where traffic congestion occurs. QMRP[CNS00] is a protocol that exhibits a very good tradeoff between the routing overhead and the success probability. In addition, QMRP has many of the good merits mentioned above. However, QMRP suffers from two drawbacks. First, QMRP deposits temporary state in the network routers for each join request. This hurts its scalability. The second drawback is that QMRP is suitable only for application with non-additive QoS requirements such as bandwidth and buffer space, and cannot be used for additive/multiplicative requirements such as delay or packet loss. While spanning join [CC97] and QoSMIC [FBP98] do no have the above problem, they have much higher overhead and lower success probability [CNS00].

This work suggests a new stateless QoS multicast routing protocol, S-QMRP, that is based on similar ideas in selecting the sub-graph to be searched for a feasible route, but uses a completely novel design. The key difference is that S-QMRP searches for a feasible route starting from the current tree towards the leaf. This has several advantages over QMRP which searches in the reverse direction from the joining leaf to the tree. First one can dynamically aggreagate several join requests and maintain a single state per multicast group. This reduces the state in the routers especially in peak hours when many join requests arrive simultaneously. It also reduces the signaling load. Second, the search in the reverse direction enables us to handle additive QoS requirements by incorporating heuristics that relay on the fact that the length of the path to the joining leaf can be obtained from the unicast routing protocol.

The new protocol maintains other merits of QMRP such as:

Robustness: The routing process is fully decentralized, which helps spreading the workload and reduces the chance of single-point falilure.

Operability: Like PIM, the protocol can operate on top of any existing unicast routing algorithm. It does not require any extra global network state to be maintained in the network.

Loop freedom: The protocol always constructs loop-free multicast trees.

We will present its supberb performance in comparison to other existing solutions, such as QosMIC and spanning join.

Joint work with Shigang Chen, Cisco Systems.

Blocking Probability Estimates in a Partitioned Sector TDMA System
Lisa Zhang, Lucent Technologies, Bell Laboratories

We consider a Time Division Multiple Access (TDMA) system in which each sector is allocated a fixed set of frequency bands. Due to spatial variation in the nature of interference, a voice call originating in a sector may have an acceptable signal to noise ratio only on a subset of the sector's frequencies. Given a Poisson traffic arrival distribution, we derive estimates of the sector blocking probability. In particular, we present lower bounds that are independent of the specific call allocation algorithm, and also examine specific call assignment algorithms such as Repacking, Least Busy and Optimal Random Routing. We develop a simple heuristic that allows us to obtain fast and accurate sector blocking probability estimates. The latter serve as crucial inputs for the design and optimization of capacity and coverage in TDMA networks.

This is joint work with Chandra Chekuri, Kavita Ramanan and Phil Whiting.

3 Future Directions and Challenges

As the Internet becomes a commercial infrastructure for supporting diverse applications, the need for service differentiation becomes important. Many of these applications are carried over several networks running different protocols and operated by competing service providers. The glue in the explosive growth of the heterogeneous internet is the simplicity of IP. However, IP-based networking is designed for delivering data traffic with best-effort (i.e., "no guarantee") service; thus it is not capable of providing end-to-end QoS.

In the last several years, many improvements have been proposed to enhance the basic Internet protocols with new functionalities to provide QoS. These functions include signaling, congestion, and admission control protocols, routing and scheduling algorithms for traffic management. Division of these functionalities between edge and core nodes identifies two fundamentally different perspectives to the QoS problem in the Internet: (i) "Integrated Services" (Intserv), and (ii) "Differentiated services" (Diffserv). Intserv offers service models, protocols, and algorithms to support end-to-end QoS on a "per-flow basis"; however, it suffers from a scalability problem due to the tremendous state information and overhead. In contrast, Diffserv proposes service models and algorithms to support QoS for "traffic aggregates" but suffers from non-deterministic and non-quantifiable QoS. There are other proposals that reside somewhere between the Diffserv and the Intserv models.

Thus, several problems emerge for future research in providing QoS, among which are: (i) managing heterogeneity, (ii) scalability, (iii) fairness, and (iv) efficient resource management.

Preserving QoS requirements of an application from one networking technology to another is a non-trivial task and currently there are no solutions for it. For example, consider an application that is initiated from a network with Diffserv support and which then traverses two networks: one with Frame Relay and the other with native ATM technology to reach its peer at another IP-Diffserv network. The initial characterization of the QoS parameters and their support over such heterogeneous networking will be significantly lost. A possible remedy is to design and operate a "virtual network" that would hide the underlying heterogeneity. Technologies such as MPLS are promising developments in this direction; however, the problem of translating QoS parameters remains.

As the Intserv approach demonstrated, QoS support per application within the network generates a large overhead in terms of computation and state information as the number of applications increases. However, if the network manages and supports traffic aggregations (a.k.a. service classes), then it is impossible to guarantee the same minimum service for all applications. It is non-trivial to define how many service classes there will be, as well as the characterization for each class, to determine how to pack the applications into these classes, and how to allocate the network resources among them. Furthermore there are fairness issues to be addressed. In addition, there is a tradeoff between the service offered and the cost to the network of the service. It is desirable to allocate the resources in a way that will (i) be as cost effective as possible, and (ii) utilize the network resources in an intelligent way.

Admission control and resource management are tightly coupled. As the emulation of circuit-switching over packet-switched networks becomes attractive, inefficient resource management may cause "service blocking". Currently, admission control is based on centralized processing, and it may be more efficient to devise distributed algorithms to support new paradigms such as "service level agreements" offered at edge routers. Given the current ownership scenarios of networks, it is a challenge to provide a flexible enough scheme for service level agreements that will allow the routing of packets through heterogeneous networks of different ownership, and still provide the desirable end-to-end service.

As the network traffic ceases to be uniform in terms of the service expectations, this brings up other issues. One such example is scheduling, where each router now has to take into account the service level of a packet before deciding to forward, keep, or drop it, within some fairness parameters (for instance, it may be necessary to keep flows from starving when the network is congested). Another example is IP-lookup, where, in a distributed routing scenario, the lookup information for a packet will now include information about its service level in addition to the usual parameters. In this scenario, it is also important that the lookups be performed extremely efficiently, especially for the high-priority packets.

In summary, the QoS problem has multiple facets and tradeoffs among conflicting objectives. Thus, a one-size-fits-all solution may not be possible, and a large part of the goal is to fit individual solutions consistently together.