

RSVP Reservation Gaps: Problems and Solutions

Kalaiarul Dharmalingam, Karol Kowalik and Martin Collier
Research Institute in Networks and Communications Engineering (RINCE)
Dublin City University, Republic of Ireland
E-mail: {arul, kowalikk, collierm}@eeng.dcu.ie

Abstract—High-end networking applications such as e-commerce, multimedia, distributed data analysis and advanced collaborative environments feature demanding end-to-end quality of service (QoS) requirements. Due to the heterogeneity exhibited by the Internet, a route from source to destination for such a flow may not be available which is comprised exclusively of QoS supporting path segments. Hence the flow must traverse one or more non-QoS path segments referred to here as *reservation gaps*. In this paper we study the problem of reservation gaps and their impact on QoS and present a solution to address the deficiencies caused by such gaps, using an Active Network approach based on the mobile agent paradigm. Furthermore, to improve the reliability in path selection and to minimise the influence of reservation gaps along the path of a QoS flow, we propose two routing algorithms, the *most reliable – shortest path (MR-S)* algorithm and the *shortest – most reliable path (S-MR)* algorithm, that select paths with the minimum number of reservation gaps. The Active Network based solution we propose works autonomously and scales to large networks such as the Internet. We demonstrate the advantages of such a solution using simulations which compares operational characteristics of QoS flows when traversing non-managed and actively managed reservation gaps. We also demonstrate the benefits of employing a routing algorithm such as *MR-S* or *S-MR* that accounts for reservation gaps in place of conventional Shortest-Path routing algorithms.

I. INTRODUCTION

As the range of applications which run on the Internet continues to widen there is a pressing need to provide deterministic service levels within the network. This requires the introduction of Quality of Service (QoS) mechanisms in the Internet. QoS refers to the capability of a network to provide priority processing including dedicated bandwidth, controlled jitter and latency, and improved loss characteristics to selected traffic classes (audio, video, etc.). IntServ/RSVP, DiffServ, MPLS, and Constraint-based routing [1] are some of the major approaches proposed to retrofit the Internet with QoS capabilities. These enhancements to the Internet are intended to provide end-to-end QoS support.

The Internet is a heterogeneous network environment interconnecting different autonomous network systems on a global scale. Due to this, the future availability of QoS support features at all nodes in the Internet is highly unlikely. Hence, non-QoS nodes will coexist with QoS-supporting nodes in the network. Throughout our discussion we will refer to the latter as “Q-nodes” and the path segments which interconnect them as Q-segments. We call the flows requiring QoS guarantees “Q-flows”. If a mixture of both Q and non-Q segments is present along a Q-flow’s path, no global end-to-end service levels can be guaranteed. We refer to the non-Q segments present along a

Q-flow’s path as *reservation gaps*. Mechanisms to compensate for the possible impact on QoS of these reservation gaps will be required if QoS-sensitive applications are to be deployed as widely as possible in the Internet.

A solution to overcome deficient reservations was presented in [2] using a receiver-initiated agent-based approach. In this paper, we present another solution using Active Networks [3] that autonomously works to manage reservation gaps present along a Q-flow’s reservation path. To permit Q-flows to traverse non-Q segments we use Q-nodes running Active Network services to monitor and provide information about the availability of resources across the reservation gap. We then manage the Q-flows traversing the gap. To increase the reliability of path selection, we propose two new routing algorithms (*MR-S* and *S-MR*) that select routes for Q-flows with the minimum number of reservation gaps. These routing enhancements are employed only at the Q-nodes. Such an Active Network approach supported by routing enhancements aims to minimise the influence of reservation gaps on the Q-flows traversing them.

The rest of the paper is organised as follows. The following section describes the reference architecture used and how reservation gaps occur therein. Section III describes the existing work on this subject and argues for an Active Network environment to solve the problem of gaps. In section IV we present the details of such a solution. In section V we describe two routing algorithms that select paths with the minimal number of reservation gaps for Q-flows. Section VI presents simulation results and section VI-C concludes the article.

II. RESERVATION GAPS

For the purpose of illustration, we describe reservation gaps in the context of the RSVP protocol [4]. The IntServ over DiffServ framework [5] provides a scalable end-to-end QoS model for the Internet. In this architecture, the stub network domains are based on an IntServ network model while the core network follows a DiffServ [6] based architecture. This approach is currently one of the most valuable solutions for end-to-end QoS provisioning, since it tries to combine the benefits of both the IntServ [4] and the DiffServ [6] architectures. This model provides QoS signalling capabilities for resource reservation by end-user applications and also provides good scalability properties when working in the core network. The reference architecture (see Fig. 1) which we have used to describe our approach for enabling robust QoS support in the Internet is based on this end-to-end QoS model.

Due to the heterogeneity exhibited by the Internet, a route from source to destination for a Q-flow may not be available

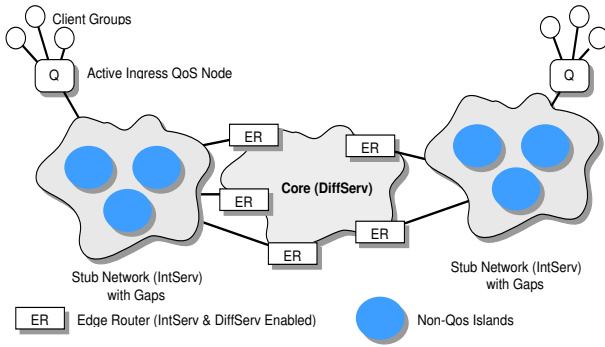


Fig. 1: IntServ Over DiffServ QoS Model With Non-QoS Islands

which is comprised exclusively of QoS supporting path segments. Hence the flow must traverse one or more reservation gaps. QoS signalling protocols like RSVP [4] for IntServ networks provide support for operation in heterogeneous networks. When non-Q-nodes (i.e. non-RSVP nodes) are present along a Q-flow's path, the RSVP signalling messages reserve required resources at the RSVP nodes (Q-nodes) and rely on the available best-effort service offered across the non-Q-nodes. Our focus in this paper is on the reservation gaps caused by the non-QoS islands in the stub network domains (see Fig. 1).

In an environment with reservation gaps, it would be possible to provide end-to-end QoS support either by: (i) restricting Q-flows to paths comprising exclusively Q-nodes; or (ii) permitting Q-flows to traverse reservation gaps assuming best-effort service provides adequate QoS across non-Q segments. However, with a low population of Q-nodes, the first approach will fail. The second approach will fail during congestion periods. A more realistic solution would be to minimise the number of reservation gaps present along a Q-flow's path and further perform monitoring of the unavoidable gaps so as to provide information about the availability of resources across the gap.

III. ACTIVE Q-NODES

A solution to overcome deficient reservations was presented in [2] using a receiver-initiated agent-based approach. In that approach, individual receiver end applications initiated mobile agents for solving the problem of deficient reservations, when reservation gaps (referred to as reservation tunnels in [2]) were detected along the end-to-end path of a Q-flow. The tunnel detection mechanism used was to back-trace and probe every node on the session's end-to-end path to identify the existence of tunnel segments. On identifying the exact location of the tunnel(s), mobile agents were deployed by the end application to monitor the tunnel characteristics and to report the measurements to the application. Our approach is scalable and is distinguished by its capacity to continuously monitor and adapt to network conditions.

Our solution requires additional features (to support management and monitoring of gaps) to be present at the Q-nodes in the network, specifically at those Q-nodes that sandwich gaps. We advocate the Active Network approach for this purpose.

The reasons for using an Active Network approach are: (i) a Q-node will not act as an entry/exit node for a reservation gap on a semi-permanent basis, but rather on a dynamic basis, as dictated by the changing network state and the operation of its routing protocols. Thus it is advantageous to implement the necessary functions using dynamically loadable modules; (ii) The start and termination point of reservation gaps can only occur at the boundaries between QoS and best-effort regions of the network; as QoS support is rolled out, the locations of these boundaries will change as will the Q-nodes which must support reservation gaps. Manual deployment of the necessary software to handle the dynamically formed reservation gaps creates a management problem which is avoided in the Active Network architecture.

We use the Netlets architecture [7], [8] for providing autonomous Active Network services. This supports dynamic loading of network services as and when required. The service code in the Netlets architecture is mobile and autonomous which eliminates the need of manual intervention for service deployment. In our discussion below, we assume that Q-nodes (supporting RSVP) in the Internet are also able to support Netlet services [7], [8]. This is a reasonable assumption, since the technology to support QoS in networks is evolving and the software (e.g., QoS signalling) on such nodes is likely to be subject to regular upgrades, as part of which Netlet [7], [8] support may be added. To support monitoring of non-Q-segments, we assume all nodes in the network support SNMP [9].

Note that the applications of the Netlets architecture are not confined to the problem of reservation gaps. Hence the justification for installing Netlet support at Q-nodes is not just to manage reservation gaps, but to host other mobile agent based solutions.

Application Level Support

For the sake of generality, we make our approach independent of the end application's in-built QoS features. We assume QoS support to non-QoS aware applications are provided by an active gateway node connecting the user to the Internet [10], such as node Q in Fig. 1. Thus a general assumption we make is that ingress nodes connecting the users to the stub network are QoS provisioned.

IV. ROBUST RESERVATION SUPPORT

Our approach to providing robust support for reservations in the presence of reservation gaps is based on the following three mechanisms: (a) discovering the reservation gaps; (b) monitoring each gap; (c) managing the Q-flows traversing the gap.

Fig. 2 depicts reservation gaps caused by a non-QoS island, G, in an IntServ network. For the example, we consider a sender node (S1) generating RSVP PATH messages destined to a receiver node. In this case, nodes Q1 and Q3 that sandwich the reservation gap Gap1 (caused by non-QoS island G) are referred to as the entry and exit nodes of the gap. In the case of a completely non-QoS provisioned stub network, the maximum path length of a reservation-gap will span from the




Fig. 2: The Reservation Gap in an IntServ Network

Ingress/Gateway QoS Node (e.g. node Q) to an Edge Node (e.g. ER) of the DiffServ domain, which is always both RSVP and DiffServ enabled.

A. Dynamically Discovering Reservation-Gaps

We use the PATH messages of the RSVP protocol to discover the reservation gaps present along a Q-flow's path. Each PATH message in the RSVP protocol includes the address of the last known RSVP-capable node in the Phop (previous hop) field. If the node receiving the PATH message does not have direct connectivity to the Phop node (based on the information available in the neighbourhood table¹), it recognises the existence of a reservation gap between itself and the node identified as the Phop node. For example in Fig 2, when a downstream node such as Q3 receives a PATH message from an upstream sender node, S1, node Q1 is identified as the Phop node, thus identifying the presence of a gap.

B. Monitoring the Reservation-Gap

On discovering the existence of a reservation gap, the relevant exit node (e.g. Q3 for Gap1 in Fig. 2), takes on the role of managing the gap and installs autonomous active service code [7], [8] at both entry and exit Q-nodes as required. Those autonomous code modules present at the entry and exit Q-nodes of the gap are referred to as *Entry-active-service* and *Exit-active-service* respectively. These services co-operate to perform monitoring of the reservation gap (for example, the services installed at Q1 & Q3 co-operate to monitor Gap1, in Fig. 2).

SNMP [9] is used to measure the path characteristics of the gap. The *Entry-active-service*, generates SNMP agents with destination *Exit-active-service*. The SNMP agents traverse the reservation gap collecting relevant metrics (for example, the available bandwidth, queue length, discard rates, interface utilisation etc.) by probing the relevant MIB entries of the non Q-nodes. The delay in traversing the gap is measured by sending time stamped packets from the *Entry-active-service* to the *Exit-active-service*. The direction of travel of the SNMP messages and delay measurement packets conforms with the travel direction of the PATH message along the reservation-gap (for example from Q1 to Q3 for a PATH message from S1 in Fig.2). Continuous packet processing will be necessary, for as long as there are non-Q-flows present in the reservation-gap.

¹A table listing the nodes to which a host node is directly connected.

C. Managing the Q-Flows Traversing the Gap

Managing the Q-flows involves intercepting and updating the RSVP signalling messages (PATH and RESV) traversing the reservation gaps to reflect the available resources in the gap. The PATH message primarily functions to install reverse routing state in each router along the path, and secondly to provide receivers with information about the characteristics of the sender traffic and end-to-end path so as that they can make appropriate reservations. The RESV messages in turn carry reservation requests to reserve resources based on the PATH message content. In the conventional reservation scheme, the PATH and RESV message are ignored along the uncontrolled reservation-gaps. In contrast, by monitoring the gap we provide the receiver and sender nodes with accurate information about the path characteristics and reservation availability.

Q-flow management across the reservation gap is accomplished as follows. The end host operating over the end-to-end QoS model (e.g. node S1 or Q in Fig. 2) sends PATH messages destined to the receiver node with a request for a reservation (Step 1 in Fig. 2). The Q-nodes (RSVP) along the path create path-state information at the local node for each PATH message and update the ADSpec object (this advertises the path characteristics to the receiver). When a PATH message is received at an exit node (Q3) of a reservation-gap, this node updates the PATH message with information gathered using the monitoring scheme (Step 2 in Fig. 2). Hence, the receiver node receives accurate state information regarding the end-to-end path. When the exit node receives the request for resource reservation (the RESV message) it checks for the availability of resources (e.g. bandwidth) and conformance to the delay constraints specified. On acceptance of the reservation by the exit node, this node informs the entry Q-node to send all data-flows belonging to the request across the exit node until the PATH state information of the flow present at the entry Q-node times out (Step 4). This is done to mimic the RSVP behaviour followed in maintaining soft-state routes along the end-to-end path. This will allow the data packets of a flow to follow the same path across the gap as that of the delay measurement packets and the RSVP signalling messages. The entry node sends data flows to the prescribed exit node by setting the Source Route Option present in the IP header of packets belonging to the flow. For example, Q-flows originating from S1 and whose signalling messages traverse Gap1 have the Source Route Option set to Q3. When the data packets reach the exit node, the Source Route Option is cleared and the packets are forwarded. This allows multiple gaps present along the path of a Q-flow to use the same Source Route Option field.

V. ROUTING ALGORITHMS TO MINIMISE THE NUMBER OF GAPS ALONG A Q-FLOW'S PATH

The Active Network based approach as described above can provide end-to-end QoS support model in a heterogeneous network environment with reservation gaps. In today's Internet RIP [11] and OSPF [12] are the two most widely used Interior Gateway Protocols (IGP). Both compute the shortest path (SP)




Fig. 3: The ISP topology

between source and destination. When working in a heterogeneous environment, such as the Internet, the shortest path may consist of an arbitrary number of both “Q” and non-Q nodes. However, it will be more efficient if the routing mechanism select paths for Q-flows with the minimal number of reservation gaps. We propose two route selection algorithms that aim to select paths with the maximum number of Q-nodes:

most reliable – shortest path algorithm (MR-S) – this selects a set of minimum hop count paths and, where there is more than one such path, selects the one with the maximum number of Q-nodes. If there are several such *most reliable – shortest* paths, random selection is used.

shortest – most reliable path algorithm (S-MR) – this selects a set maximum number of Q-nodes paths and, where there is more than one such path, selects the one with the minimum hop count. If there are several such *shortest – most reliable* paths, random selection is used.

We assume here that non-Q nodes forward packets according to decisions of existing routing protocols (such as RIP or OSPF) while the Q-nodes use the MR-S or S-MR routing algorithms.

VI. SIMULATIONS

We present two distinct sets of simulations to study the problem of reservation gaps. In the first set, we present a comparative analysis between the most reliable – shortest path (MR-S), the shortest – most reliable (S-MR) and the traditional shortest path algorithms. In the second set, we compare operational characteristics of Q-flows when traversing non-managed and actively managed reservation gaps.

A. Routing Enhancements

1) *Network Model*: We evaluated the performance of three routing approaches (SP, MR-S and S-MR) described in Section V with the so-called ISP topology [13] shown in Figure 3. In general, the network topology is assumed to consist of N nodes connected using L bidirectional links each with capacity C (for the ISP topology we have used $N = 18$, $L = 30$ and $C = 20$). We also assume that within the network there are N^Q Q-nodes and N^{nQ} non-Q nodes (where: $N^Q + N^{nQ} = N$).

The requests arrive at each node independently according to a Poisson distribution with rate λ and have exponentially distributed holding times with mean value $1/\mu$. The requested amount of bandwidth is uniformly distributed over the interval:

[64kb/s, 6Mb/s], with mean value $B = 3.32\text{Mb/s}$. If traffic is generated by N^s source nodes, it produces the network offered load [14]: $\rho = \lambda N^s B h' / \mu L C$, where h' is the average shortest path distance between nodes, calculated over all source-destination pairs (for the ISP topology: $h' = 2.36$ if $N^s = 18$). In our experiment we have used a mean connection holding time of 60sec and choose λ to produce the required offered load in the network.

2) *Performance metrics*: In our simulations we assumed that, when a new request arrives it can receive one of two responses if path monitoring is used:

- accept - if there are enough resources along the chosen and monitored path;
- reject - if resources along the chosen and monitored path cannot accommodate the new request.

However if path monitoring is not present (as in the existing Internet) a third response can be generated:

- fail - if the decision was to accept a connection on the path, but the path failed to provide the required QoS level. This occurs if no information is given about the lack of resources (due to the absence of monitoring) and we assume that the connection failed due to a user decision, since the granted QoS level does not conform to the requested quality.

We have assumed that the lifetime of failed connections is exponentially distributed with a mean value equal to half of the mean value of a standard connection ($1/2\mu$). We will explain why such a value was chosen in Section VI-A.3.

We have used the following metrics to compare the performance of the three routing approaches (SP, MR-S, S-MR): the *call blocking rate* – defined as:

$$\text{call blocking rate} = \frac{\text{number of (rejected + failed) requests}}{\text{number of arrived requests}}$$

which is used to calculate the probability of rejecting the new request; the *average path length* – defined as:

$$\text{avg. path length} = \frac{\sum \text{path length of accepted connections}}{\text{number of accepted connections}}$$

used as an indicator of resource consumption when comparing the algorithm to algorithms which limit the hop count; the *installation cost* – defined as:

$$\text{installation cost} = \frac{\text{number of monitoring points}}{\text{number of monitored connections}}$$

which is used to calculate the cost of installing active services along the path; and the *reliability* – defined as:

$$\text{reliability} = \frac{\text{number of QoS aware nodes along the path}}{\text{number of total nodes along the path}}$$

(a path is assumed to be more reliable if it has a higher ratio of Q-nodes).

3) *Results*: The two proposed routing algorithms (MR-S and S-MR) aim to improve the reliability of routing protocols in a heterogeneous environment. Hence for the given network with the ISP topology we have randomly increased the number of Q-nodes starting from a network which does not provide any QoS support ($N^Q = 0$ and $N^{nQ} = N$) until we have reached




Fig. 4: Call blocking probability of SP, MR-S and S-MR under increasing number of QoS aware nodes




Fig. 5: Average length of the path chosen by SP, MR-S and S-MR

the fully QoS supportive network ($N^Q = N$ and $N^{nQ} = 0$) – we call this a cycle of simulation. In our experiment each cycle of simulation was repeated 500 times and we present the average results below. As shown in Fig. 4 when path monitoring is not supported and the shortest path (SP) is chosen, the blocking probability is quite high for networks with a small number of Q-nodes. This is caused by connections that are setup although there are not enough resources to accommodate them and which fail after establishment. This type of failed connection uses resources unnecessarily and so blocks other potential connections. We can only reduce the blocking probability when using SP by increasing the number of Q-nodes. We used $1/a\mu$ as the mean holding time of failed connections. We assumed that $a > 1$ because we expect that a failed connections will be terminated earlier than had it been successful. In our simulations we used $a = 2$. Using other values of a , which are greater than one, changes the blocking probability but this remains a monotonic decreasing function of the number of Q-nodes.

If a path monitoring mechanism is employed in conjunction with routing algorithms which select more reliable paths (as in the MR-S and S-MR curves in Figure 4) the blocking probability is reduced. By using monitoring we prevent situations arising where, due to a lack of accurate reservation information, connections are established over links with insufficient resources. The MR-S algorithm, which chooses the most reliable path from the set of shortest paths, shows the extent to which blocking probability can be reduced by path monitoring. The S-MR algorithm also reduces blocking probability, when the number of Q-nodes is small. However when the number of Q-nodes is more than half of the total number of nodes, it generates a high blocking rate, because of the use of non shortest paths which consume extra resources. This is clearly seen by comparing Figures 5 and 4. This also confirms the findings of other researchers [14], that algorithms limiting hop count produce a lower blocking probability. The average path length of SP (the bottom curve in Figure 5) grows slightly when the number of Q-nodes increases. Clearly, when information about resources available in the network is not provided, setup of connections requiring only a few hops is more likely to be




Fig. 6: Reliability of routing decisions of SP, MR-S and S-MR

successful than of connections where the source and destination nodes are further apart.

Although S-MR produces a higher blocking probability than MR-S, it is more reliable when compared with the SP and MR-S algorithms (see Figure 6). This suggest that the S-MR algorithm should be used only for connections requiring high reliability and producing higher revenue. Other flows should be processed using MR-S.

When evaluating the cost of installing the monitoring mechanisms in Figure 7 we can see that the cost of using monitoring for S-MR is comparable with that for MR-S (when the number of Q-nodes is less than half of the total number of nodes) or is even lower (if the number of Q-node is greater than half the total number of nodes). So if a network administrator decides that the computational cost of monitoring non-Q segments is excessive, he could use S-MR even if it produces a higher blocking probability than MR-S.

B. Non-Managed Vs Actively Managed Gap

We expect that, in a heterogenous environment like the Internet, a minimal number of reservation gaps will occur even when using routing algorithms such as MR-S and SR-M to select paths for the Q-flows. Hence it is desirable to study the characteristics of Q-flows when traversing managed and actively managed reservation gaps. For the example, we




Fig. 7: Cost of path monitoring




Fig. 9: Actively Managed Reservation Gap




Fig. 8: Service Degradation For QoS Flows across a Gap

considered three UDP based Q-flows traversing a non-QoS link, L, of capacity 1Mb/sec. The UDP traffic generator is a constant bit rate (CBR) source with exponentially distributed on and off periods. We define Q-Flow 1, and Q-Flow 2 with bandwidth requirement $\approx 0.4\text{Mb/sec}$ and 0.9Mb/sec respectively. When a new Q-flow, Q-Flow 3, with bandwidth requirement $\approx 1.5\text{Mb/sec}$, greater than the available resources, enters the reservation gap, a very high percentage of the link's capacity is absorbed by this flow (Q-Flow 3). This causes degradation to the existing Q-flows. This is shown in Fig. 8. Overall, all the flows suffer heavy packet losses and the network resources are inefficiently utilised.

With a robust QoS support for reservation, the third flow, Q-Flow 3, is informed about the non-availability of resources across the non-QoS link, L (in Fig. 9, at $t \approx 300$, the third flow makes a setup request but is rejected). Hence it does not interfere with the existing Q-flows (Q-Flow 1, Q-Flow 2) which continue to receive the requested QoS (see Fig. 9) and the network resources are efficiently utilised.

C. Conclusions

The unpredictable behaviour of traffic within the non-QoS path segments present along a Q-flow's path and the inability to support reservations across them can cause problems in providing end-to-end service guarantees in the Internet. We have described an Active Network solution using the mobile agent

paradigm to build a robust end-to-end QoS support model. We also proposed routing enhancements (*MR-S* and *S-MR*) that when employed at Q-nodes select a path for the Q-flow with the minimum number of reservation gaps. Our technique features excellent dynamics and scales for large networks and user populations. The good dynamics make support of short lived Q-flows feasible. The control traffic generated (to monitor and manage the non-QoS path segments) is confined to the corresponding reservation gap, thus reducing congestion and packet loss. The technique described here makes it possible to deploy applications in the network which have quite hard QoS guarantee requirements, even when a significant number of network nodes support only best-effort service. Such techniques will be of critical importance in ensuring the graceful transition of the Internet from a best-effort service model to a service model featuring QoS guarantees.

REFERENCES

- [1] X. Xiao and L. M. Ni, "Internet QoS: A Big Picture," *IEEE Network*, vol. Vol.13, Mar 1999.
- [2] H. de Meer, J.-P. Richter, A. Puliafito, and O. Tomarchio, "Tunnel agents for enhanced Internet QoS," *IEEE Concurrency*, vol. 6, no. 2, pp. 30–39, 1998.
- [3] K. Psounis, "Active Networks: Applications, Security, Safety, and Architectures," *IEEE Communications Surveys*, First Quarter 1999.
- [4] L. Zhang, S. Deering, and D. Estrin, "RSVP: A New Resource ReSer-Vation Protocol," *IEEE network*, vol. 7, September 1993.
- [5] Y. B. et. al, "A Framework for Integrated Services Operation over DiffServ Networks," *RFC 2998*, Nov.2000.
- [6] S. B. et. al, "An Architecture for Differentiated Services," *RFC 2475*, Dec.1998.
- [7] M. Collier, "Netlets: The Future of Networking?" *IEEE OpenArch*, April 1998.
- [8] K. Dharmalingam and M. Collier, "Netlets: A New Active Network Architecture," *IEE/IEEE Symposium on Telecommunications Systems Research, Dublin, Ireland*, Nov 2001.
- [9] W. Stallings, "SNMP, SNMPv2 and CMIP. The Practical Guide to Network Management Standards," *Addison Wesley*, 1993.
- [10] Kalaialur Dharmalingam and Martin Collier, "Transparent QoS Support For Network Applications using Netlets," *Proceedings of Mobile Agents for Telecommunication Applications (MATA)*, Spain, Oct 2002.
- [11] G. Malkin, "RIP Version 2," *RFC 2453*, November 1998.
- [12] J. Moy, "OSPF Version 2," *RFC 2178*, April 1998.
- [13] G. Apostopoulos, R. Guerin, S. Kamat, A. Orda, and S. K. Tripathi, "Intradomain QoS Routing in IP Networks: A Feasibility and Cost/Benefit Analysis," *IEEE Network*, 13(5):42–54, Sept./Oct. 1999.
- [14] A. Shaikh, J. Rexford, and K. S. Shin, "Evaluating the impact of stale link state on quality-of-service routing," *IEEE/ACM Transactions on Networking*, April 2001.