

Quality of Service (QoS)-Based Routing in the Internet - Some Issues

Status of this Memo

This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress."

To learn the current status of any Internet Draft, please check the `1id-abstracts.txt` listing contained in the Internet Drafts Shadow Directories on `ds.internic.net`, `nic.nordu.net`, `ftp.nisc.sri.com` or `munni.oz.au`.

Distribution of this memo is unlimited.

This Internet Draft expires December, 16, 1996.

Abstract

There are many reasons to consider QoS-based routing as a component of the integrated services Internet. But several questions arise with regard to its development: what are the requirements on QoS-based routing in the Internet? What sort of a routing architecture is practical? What are the technical and policy issues that arise in realizing a QoS-based Internet routing architecture? This draft is an attempt to generate some discussion on these topics. To this end we present some potential requirements on path computation, efficiency, robustness and scalability, and describe some issues in realizing a QoS-based routing architecture. Our conclusions are that QoS-based routing is a challenging problem, and that it may be best to address the intradomain, interdomain and scalability aspects separately in developing the routing architecture.

1. INTRODUCTION

The internet is being used increasingly for transport of multimedia information such as image, voice and video. With the increase in sophistication of desktop computers, and the availability of networked multimedia applications, the Internet is likely to see more of this type of traffic. This type of usage, however, is ad-hoc in that the IP network architecture has been designed for "best-effort" delivery, without any guarantees on throughput or delay. It has been recognized that to adequately support realtime traffic flows with bandwidth and delay requirements, the following features should be present in the Internet [1]:

- new service classes beyond best-effort to provide certain guarantees on throughput and delay to applications;
- a mechanism that allows applications to signal to the network their quality of service (QoS) requirements; and
- traffic management mechanisms in hosts and routers

The IETF has been working on the first two issues: defining service classes in addition to best-effort [1], and a resource reservation protocol (RSVP) [2] that applications may use to reserve network resources to get the level of service they desire. Traffic management features are also being built by router vendors. An issue being discussed now is whether a routing architecture that allows the dynamic discovery of QoS-accommodating paths in the network for data flows, in the presence of changes in network topology and loading, is required. In our opinion, without QoS-based routing, the definition of new service classes and signalling protocols may only be of limited use in supporting real-time applications.

Our assumption is that it is not economical to overprovision network resources in order to obviate the need for QoS mechanisms. Under these circumstances, there is a justifiable need for QoS-based routing. First, without it, protocols like RSVP may convey the QoS requirements of a flow to routers, but there is no guarantee that an existing acceptable path between the sender and the receiver(s) will be found and utilized. An alternative, suggested often, is to let routers keep track of multiple paths and attempt flow setup by trial and error along these paths. This technique does not scale well with network diameter, and the blocking probability and setup time may be high.

Second, economic network engineering depends on the use of an efficient routing scheme. For instance, with a traditional IP routing algorithm, a single fixed path is selected for all traffic between a pair of nodes. To assure the desired QoS, this path must be engineered to accommodate the peak traffic demand between the node pair. This results in the inefficient use of resources, since there may be other underutilized paths between the same node pair. On the other hand, a QoS-based routing algorithm would allow the network capacity to be

engineered efficiently by virtue of its ability to take into account the current network state in making routing decisions.

Finally, a network state-dependent routing scheme can compensate for imprecise network engineering. Specifically, when the traffic load

exceeds the engineered limits due to exceptional circumstances (e.g., focussed overload after failures), a QoS-based routing scheme allows more traffic to be routed and a more graceful performance degradation as compared to a state-insensitive routing scheme [4]. Indeed, these are the reasons for the development of state-dependent routing schemes for long distance phone networks [5] and ATM networks [6]. QoS-based routing must therefore be considered an essential component of the overall scheme to provide QoS guarantees to applications in the Internet.

Given the need for QoS-based routing, the following questions arise: what are the requirements on QoS-based routing in the Internet? What sort of a routing architecture is practical? What are the technical and policy issues that arise in realizing a QoS-based Internet routing architecture? This draft is an attempt to generate some discussion on these topics. To this end, in the next section, we point out some potential requirements. In Sections [3](#) and [4](#), we describe some issues that arise in developing a QoS-based routing architecture satisfying these requirements. Related work is outlined in [Section 5](#), and summary and conclusions are presented in [Section 6](#).

[2.](#) SOME REQUIREMENTS ON QOS-BASED IP ROUTING

The foremost issue in developing a QoS-based IP routing architecture is the definition of the requirements it must satisfy. Given the organization of the internet, the nature of the traffic, the service classes being considered, and the requirements of the users, there could be many requirements on QoS-based routing. To make a start, we may consider the following preliminary requirements:

Metrics and Paths

- Support for multiple path metrics (bandwidth, delay, etc.). The definition of the metrics would be based on the service classes defined by the intserv WG.
- The ability to dynamically determine paths satisfying multiple constraints (e.g., bandwidth and delay).
- QoS support for multiple types of flows, i.e., unicast and many-to-many multicast.

Robustness

- The ability to detect and respond to dynamically changing QoS capabilities of a link.
- The ability to continue routing in the presence of multiple, simultaneous failures in the network.

[draft-nair-qos-based-routing-00.txt](#)

Page 3

Efficiency

- The ability to route short-lived data flows with minimal overhead.
- The ability to utilize network resources efficiently, through load spreading, global admission control techniques and the like.
- Minimization of routing overheads.
- Support for resource control, i.e., the ability to limit resource consumption by different classes of traffic.

Scalability and Priority

- The ability to scale to large numbers of nodes and links, and to a large network diameter.
- Support for prioritizing flows, and to permit high priority flows to be established with precedence over lower priority flows.

Interdomain and Policy

An area that requires some thought is the requirements for interdomain routing. Specifically, the nature of the information that is exchanged at the border of two domains must be determined. If the information exchange is compact, the interdomain routing scheme can be relatively simple ([Section 3.8](#)). In any case, the ability to introduce policy constraints on routing information flow at the boundary between two organizations is required.

Other

Other requirements, such as mobility support, may be defined.

These are complex requirements, and a number of challenging problems must be solved to satisfy them. The resulting routing scheme can be expected to be more sophisticated than the existing internet routing schemes. It is, however, encouraging that recent Internet routing work, such as Nimrod and Source Demand Routing, has addressed scalability, on-demand route computation and source routing issues. The ATM Forum PNNI work illustrates an effort to develop a scalable, QoS-based routing architecture. To what extent the ideas developed under these efforts may be incorporated in the QoS-based routing architecture for the Internet is a subject that requires much discussion. This draft does not address that question. Instead, our objective is to describe some of the problems that arise in realizing a QoS-based Internet routing architecture.

To this end, we first consider intradomain routing, without scalability requirements. As a reference routing architecture for this case, link state routing with distributed intelligence seems to be a reasonable choice. To build QoS-based routing features in this architecture, algorithm design issues, for instance, computing "low cost" multicast distribution trees for flows, or efficient broadcast

[draft-nair-qos-based-routing-00.txt](#)

Page 4

schemes for propagating QoS information must be addressed. Other problems that involve policy and scalability issues in the global, multiply administered internet may be addressed separately, once the requirements are determined. This approach allows us to focus on two distinct classes of problems that arise in realizing the QoS-based routing architecture for the Internet.

Following the requirements above, the intradomain link state architecture should allow a router determine a feasible path for a given unicast flow on demand. For multicast flows with dynamic receiver sets, QoS-based routing should allow the dynamic, incremental computation of QoS-accommodating multicast trees. Under link state routing, each router in an autonomous system (AS) monitors the QoS available on each directly attached link, and its own resource availability. This information is broadcast to other routers in the AS on a periodic and/or event-driven basis. Based on the view of the network state, a router may determine the path for a given flow. Source routing, along with local and global admission control techniques may be used to accept or reject a flow along the path. Also, routers may recognize the priorities of flows (perhaps signalled via RSVP) and implement a mechanism to ensure a high rate of success for guaranteeing QoS for high priority flows.

The issues that arise in intradomain QoS-based routing may be broadly classified as follows:

- How do routers determine the QoS capability of each outgoing link

and reserve link resources? Note that some of these links may be virtual, over ATM networks.

- How do routers propagate the QoS knowledge to remote routers to aid in path selection?
- How are QoS-accommodating paths computed by routers for unicast flows?
- How are QoS-accommodating paths computed for multicast flows?
- What is the mechanism for prioritizing flows?
- What is the mechanism for resource control?
- What are the performance impacts of dealing with many short-lived flows (e.g., web page accesses)?
- How does the RSVP model fit with the QoS-based routing architecture?

Each of these points is discussed in the next section. Scalability and interdomain routing issues are described briefly in [Section 4](#).

[3.](#) INTRADOMAIN ROUTING ISSUES

[3.1](#) QoS Determination and Resource Reservation

The integrated services working group of the IETF has concentrated on local resource management within routers. In particular, the notion of "admission control" has been introduced that allows a router to make a

[draft-nair-qos-based-routing-00.txt](#)

Page 5

decision as to whether a given flow may be accepted or rejected based on local resource availability. The admission control process requires knowledge of the QoS available on all links attached to a router. It is still an open issue, however, as to how the QoS values are computed for various link types such as multiple access links. The resolution of this issue is in the domain of the ISSLL working group.

The intricacies of this problem may be appreciated when the case of a router connected to a large non-broadcast multiple access network, such as ATM, is considered. In this case, a router may have multiple next hop choices, with corresponding paths across the ATM network, each with possibly different QoS availability. The issues are, how does a router determine what the routing options are across an ATM network, what the QoS availability over each available route is, and what QoS value to advertise for the ATM link when QoS-based routing is implemented in the wider internet. To put this problem in perspective, the currently

defined standards for IP routing over ATM, such as NHRP, allow the selection of a single exit point in the ATM network for an IP datagram. With QoS requirements on IP flows, there may not be an ATM path which accommodates the given QoS from the entry router to the exit router returned by NHRP. An approach like IPNNI [7] would be helpful here, although with some complexity. A related problem is the reservation of resources over a broadcast network, say an ethernet. Because access to the network is distributed, some coordination is required among routers in reserving resources.

3.2 Knowledge Propagation

The link state architecture requires routers to broadcast the local resource status (such as available link capacity, delay measurements, etc.) and the local topology information to all other routers in the AS. The broadcast of status information from one router to others might be an expensive process in terms of communication and processing overheads. So far, intradomain routing has been based on fixed link metrics (i.e., minimum hop, or shortest-path based on static metrics). In this environment, efficiency of information propagation has not been an issue. Thus, routing algorithms such as OSPF have used flooding with periodic refreshing as a means to propagate updates. Given that linkstate routing is essential for efficient QoS-based routing, flooding-type mechanisms are not suitable for information broadcasting. Alternative techniques to reduce broadcast overheads, such as tree-based forwarding, have been proposed [8]. These have to be implemented. In addition, link status information such as utilization can be volatile, i.e., their values can change significantly in a short period of time with the admittance of real-time flows. How should such information be quantized in order to reduce the update generation rate is an issue to be resolved. Clearly, the less accurate the information, the less likely that a feasible path will be found by a source router. On the other hand, the overheads of keeping very accurate information at each router may be high. Aggregation of status information reduces the frequency of update generation, but how it affects routing algorithm performance has to be determined.

3.3 Unicast Path Computation Algorithms

Given the availability of network state information at each router, how should paths be computed for unicast flows? The computation of a "feasible" path for a given flow is not difficult: a router just

eliminates all infeasible links and nodes, i.e., those that cannot support the requested QoS, from its local representation of the network topology and finds an end-to-end path in the remaining topology. But as studies in circuit-switched networks have shown that even in limited topologies this sort of "feasible-path routing" is unacceptable, i.e.,

it results in the admission of lesser number of flows into the network than what is possible otherwise. Moreover, as noted above, aggregation of link status information is usually lossy. This aggravates the problem of flow blocking.

Feasible-path routing corresponds to the notion of local admission control defined by the IETF's integrated services group; a flow is routed over a path as long as it passes the admission control criterion of each router along the path. While this type of admission control is required to control the subscription of resources at a router, it does not take into account any global view of resource consumption by individual flows. Efficiency in routing is achieved only when an additional layer of admission control is implemented. This higher level admission control procedure would consider the resource requirement of each flow in relation to the available resources along a path in order to determine whether it is profitable overall to admit the flow [9]. Thus, an individual flow may be rejected even if there are feasible paths to route it, if it is found that admitting the flow would result in an overall decline in the number of flows carried by the network. The formulation of this higher level admission control, with fairness to all flows desiring entry to the network, is an area that requires work.

Path computation with multiple QoS constraints on a flow is a difficult problem. Indeed, for some combinations of QoS constraints, the problem is NP-complete. However, algorithms have been proposed for computing paths with combined bandwidth and delay constraints [10], and such algorithms can be incorporated in the routing scheme.

3.4 Flow Prioritization

Given that critical flows must be accorded higher priority than other types of flows, a mechanism must be implemented in the network to recognize flow priorities. It is assumed that RSVP can be used to signal the priority of a flow. There are then two aspects to prioritizing flows. First, there must be a policy to decide how different users are allowed to set priorities for flows they originate. The network must be able to verify that a given flow is allowed to claim a priority level signalled for it. Second, the routing scheme must ensure that a path with the requested QoS will be found for a flow with a probability that increases with the priority of the flow. In other words, for a given network load, a high priority flow should be more likely to get a certain QoS from the network than a lower priority flow requesting the same QoS.

The policy and verification aspects are outside the scope of this draft. The routing mechanism for implementing flow priorities may be based on preemption combined with dynamic rerouting of lower priority flows. For example, assuming a small, fixed number of priority levels (e.g., 16), each router may broadcast the resources locally allocated to flows of each priority level. A router, when computing a path,

first attempts to find a path that has the necessary unallocated

resources to support the QoS requirements of the flow. If such a path cannot be found, the router attempts to find a path where adequate resources may be freed by displacing lower priority flows. In the end, a high priority flow may be routed by preempting resources allocated to lower priority flows. Routers may attempt to dynamically reroute preempted flows using the same procedure.

Routing procedures for flow prioritization may thus be complex. Identification and evaluation of different procedures are areas that require investigation.

[3.5](#) Resource Control

Given the existence of multiple service classes, it is necessary to engineer a network to carry the forecasted traffic demands of each class. To do this, a network administrator may logically partition router and link resources among various service classes. Strict partitioning, however, may result in inefficient use of network resources, since there may be periods of time during which there may be excessive traffic load under one class and light load under another. It is thus desirable to allow unused resources to be allotted traffic that needs them. This type of sharing, however, must be done in a controlled fashion in order to prevent traffic under some service class from taking up more resources than what was engineered for it for prolonged periods of time. This is the job of the routing scheme. This may be done via preemption, in a manner not contradictory to the priority assignment of active traffic flows. Non-preemptive dynamic resource sharing techniques are possible, as illustrated by the Real Time Network Routing architecture of the AT&T phone network [5]. The design of an appropriate resource sharing scheme, and its incorporation into the QoS-based routing scheme is a challenging issue. In particular, the overheads incurred by the routing scheme in the implementation of logical resource partitioning and dynamic sharing is an issue that must be studied carefully.

[3.6](#) Short-Lived Flows

The QoS-based routing model incurs certain overheads during flow establishment, for example, computing a source route. Whether this overhead is disproportionate compared to the length of the sessions is an issue. In general, this problem arises in virtual circuit networks such as ATM, where many short-lived SVCs result in increased call set-up overheads. Even without QoS-based routing, RSVP flow establishment overheads are incurred for each session. An area worth investigation is the minimization of routing-related overheads during flow establishment. One approach that is useful is caching recently

computed routes, so that new sessions to the same destination do not cause recomputation of paths. The issue of efficient flow set-up in general needs investigation.

[3.7](#) QoS-Based Routing for Multicast Flows

Computing QoS accommodating paths for multicast flows is a tricky problem, especially if the notion of higher level admission control ([Section 3.3](#)) is included. The dynamism in the receiver set allowed by IP multicast also adds to the problem. In any case, routers in an AS must keep track of the id of subnets where group members are present, along with the QoS requested by them, in order to compute the

[draft-nair-qos-based-routing-00.txt](#)

Page 8

appropriate multicast trees. This corresponds to an enhanced MOSPF-type algorithm [20]. As receivers leave or join a multicast group, existing trees from different sources may have to be incrementally modified. Many such heuristic incremental algorithms are presently known [23-25]. Computing optimal shared trees based on the shared reservation style [2] may require new algorithms. Finally, scalability is an issue with algorithms that require knowledge of receiver sets.

[3.8](#) RSVP and QoS-Based Routing

The QoS-based routing model has some implications on signalling. Specifically, under this routing model, the QoS desired for a flow must be specified before a path can be computed. In contrast, under RSVP, a path must exist before a reservation can be made. RSVP utilizes PATH messages to notify the receivers of a session and the traffic characteristics at the sender, and to establish the flow's path state in the routers. With QoS-based routing, the latter function is not useful since the flow path is computed based only on the receivers' reservation. Using the current RSVP model with QoS-based routing, a unicast flow establishment would require a PATH message sent from the source to the receiver along a best effort path, say, and a RESERVE message sent over a QoS-accommodating flow computed by a router close to the receiver.

With regard to multicast flows, the path for a multicast flow under the current RSVP specifications is the IP multicast tree of the source. To receive a flow, a host must first join the multicast group, and hence the multicast tree for a given source. To request a certain QoS for the flow, each receiver must send a RESERVE message towards the root of the tree, reserving link and router resources enroute. The set of receivers of a multicast flow is dynamic, and transparent to the source of the flow. The multicast routing protocol dynamically maintains a tree per source, in which the leaves are the current set of receivers. This protocol establishes the multicast tree independently of the QoS requirements of flows subsequently routed over it. Thus,

there is no guarantee that a leaf would be successful in reserving the resources to get the desired QoS.

One solution for QoS-based multicast routing is to separate RSVP PATH message flow and actual data flow from a source. Thus, a basic multicast routing protocol would provide a tree for delivering PATH messages to any receiver joining the group. The reception of a RESERVE message by a router would trigger an algorithm to (incrementally) compute a new tree for the data flow, i.e., addition/deletion of branches to the existing tree, as discussed in the previous section.

4. SCALING AND INTERDOMAIN ROUTING

4.1 Scaling by Hierarchical Aggregation

A scalable intradomain routing architecture is needed to handle the thousands of nodes that may exist within an AS. Scaling by hierarchical aggregation is a common technique, as exemplified by the ATM Forum PNNI routing scheme [12]. Hierarchical aggregation, however, introduces problems with regard to the accuracy of the aggregated state information [13]. Also, the aggregation of paths under multiple constraints is difficult. One of the difficulties is the risk of

[draft-nair-qos-based-routing-00.txt](#)

Page 9

accepting a flow based on inaccurate information, but not being able to support the QoS requirements of flow because the capabilities of the actual paths that are aggregated are not known at the source. Much study is needed to understand and characterize the performance impacts of aggregating path metric information. Meanwhile, a way to compensate for inaccuracies is to use crankback, i.e., dynamic search for alternate paths as a flow is being routed. Crankback, however, increases the time to set up a flow, and also results in overall poor utilization of resources. Thus, crankback is the mechanism of last resort, and the minimization of its use is a problem that requires work.

4.2 Interdomain Routing

Many issues arise when multiple ASs, each implementing QoS-based routing within their networks, interface. Efficient end-to-end QoS-based routing across multiple ASs require some exchange of information between them, but individual ASs may not wish to reveal too much information. For instance, it may not be possible or desirable for an AS to reveal internal topology information to others. Also, even if there are no policy constraints on exporting information, overall scalability of the internet routing architecture may not allow this. Thus, the issue is what (minimal) information should be exchanged between different autonomous systems to implement end-to-end QoS-based routing?

It is desirable that information flow across ASs should be much more abridged than what flows inside an AS. For example, fairly stable QoS information applicable to the entire network may be advertised. This means that an AS should engineer its network carefully to ensure that it can meet the QoS advertised. Compact information flow may allow the implementation of QoS-based routing through enhanced versions of traditional interdomain protocols such as BGP.

The precise definition of an interdomain routing architecture, which accommodates QoS and also aids in establishing routing policies is an area that requires work. The scalability of any proposed architecture must be evaluated, along with its efficiency in large configurations. Finally, interdomain QoS-based multicast routing is a challenging problem.

5. PROGNOSIS AND RELATED WORK

QoS-based routing in the Internet is indeed a daunting problem. To make some progress in this area, it may be necessary to make a modest beginning and concentrate on a intradomain routing protocol, say, a QoS-enhanced version of OSPF. Scalability via hierarchical aggregation, in our opinion, is a problem that can be separately tackled. Similarly, interdomain routing is also an orthogonal issue. At the intradomain level, the development of a routing scheme incorporating some of the basic requirements outlined in this draft is achievable. Indeed, commercially available proprietary routing solutions for frame relay and ATM networks from vendors such as Cascade, FORE, Stratacom and others demonstrate many of the desirable QoS-based routing features. The following is a summary of related work in various areas, and the missing pieces.

5.1 IP Routing Schemes

Among intradomain routing schemes, OSPF [3] is a link state routing protocol, utilizing distributed state information. Under OSPF, network topology, as well as an administratively configurable metric for each link are propagated throughout the network using flooding. Each node has its own view of the network state, and a shortest path algorithm is used to compute the destination-based routing table. OSPF allows the definition of a two-level routing hierarchy, and hence a good degree of scaling. As compared to OSPF, a link state QoS-based routing requires the dynamic distribution of link and nodal state information, as the corresponding resource availability changes. For scaling purposes, this would require a tree-based update propagation scheme instead of flooding. Also, instead of destination-based routing tables, a QoS-based routing scheme would use on-demand path

computation during flow establishment. This places an overhead on routers, and efficient schemes have to be designed to minimize this overhead. Finally, source routing, at least for flow establishment, has to be used instead of destination oriented next hop forwarding. Some state information will also be maintained by routers about each flow.

Interdomain routing protocols, such as BGP [14], primarily focus on the control of information flow between administratively distinct routing domains. As such, they do not maintain fine state information about routing domains. Rather, topological reachability of domains, subject to user-defined preferences for alternate route selection, is the goal. Thus, BGP uses an enhanced distance-vector routing scheme (called, path vector [15]), with destination-oriented hop by hop forwarding. QoS-enhanced BGP requires work.

Among the newer approaches, Source Demand Routing (SDR) [16] allows on-demand path computation by routers and the implementation of strict and loose source source routing. The SDR approach can be used for both intradomain and interdomain source routing. The Nimrod architecture [17] has a number of concepts built in to handle scalability and specialized path computation. These are applicable in addressing the relevant aspects of QoS-based routing.

[5.2](#) Internet Multicasting

IP multicasting is mainly concerned with establishing multicast trees subject to changes in topology and receiver group membership [18]. IP multicast schemes utilize the underlying unicast routing protocols to establish the multicast trees. For example, the Multicast OSPF (MOSPF) [19] works with OSPF. As with QoS-enhanced OSPF, it is possible to consider a QoS-accommodating MOSPF.

More recent multicast routing protocols, such as Core-Based Trees [20], and Protocol-Independent Multicast [21] do not rely on specific types of unicast routing schemes. They may need extensive revision to accommodate QoS. The MOSPF mode of operation is in fact more suitable for QoS-based multicast routing, since it is possible to tap the networkwide QoS status from the underlying link state routing scheme. Given the availability of QoS information, many heuristic algorithms to compute multicast trees are known [22-24]. The incorporation of flow aggregation and the sharing of multicast trees by several sources, as defined under RS-VP, into these algorithms requires work.

[5.3](#) ATM Routing

The ATM PNNI routing scheme [12] is a hierarchical link state QoS-based routing scheme. It incorporates mechanisms for QoS-based path computation and scalability, but it does not address some of the

problems outlined earlier: routing of many-to-many multicast flows of the type required by RSVP; interdomain policy routing issues; QoS measurements on shared links; global admission control; flow prioritization; efficient broadcast of state information; and performance issues.

There has recently been some work on SVC routing schemes for ATM networks. This research illustrates various approaches to achieving higher throughput from routing schemes, as compared to single path shortest-path routing. The flows considered usually have a single QoS requirement, i.e., bandwidth. The following is a sampling of work in this area. In [25], Sibal and Desimone describe a routing strategy for multihop networks, modelled after the reservation-based alternate routing strategy for two-hop phone networks [26]. This generalized algorithm is based on distance vector routing, and allows only a single bandwidth value for all flows. It also assumes that the input traffic distribution is known apriori. In [27], Bahk and Zarki analyze the throughput performance of several heuristic algorithms for QoS-based routing. These algorithms are simple variants of shortest-path, and some of them require coarse link state information. Mitra, Morrison and Ramakrishnan present a scheme for optimal network design [6] that assumes centralized routing and a small number of bandwidth classes. Also, input traffic is assumed to have certain characteristics. Gawlick, Kalmanek and Ramakrishnan present an on-demand routing scheme for PVCs [28] and show its superiority over shortestpath. This algorithm too is centralized. Finally, Ahmadi, Chen and Guerin present a routing and admission control heuristic algorithm, again with better performance than shortest-path [9].

The contributions of prior research in this area is of significant importance, lending some insights into unicast path computation procedures. The practical aspects of routing, such as overhead reduction, and efficiency under a mix of traffic and information aggregation, scalability, etc., have not been investigated thoroughly. The definition of a practical IP QoS-based routing architecture, and its implementation, requires new work as outlined in this draft.

6. SUMMARY AND CONCLUSIONS

There are many reasons to consider QoS-based routing as an essential component of the overall integrated services Internet infrastructure. The first step in developing a QoS-based routing architecture is the specification of its requirements. In particular, the requirements for interdomain routing need much discussion. In this draft, we presented some potential requirements on path computation, efficiency, robustness and scalability, and described some issues in realizing a QoS-based routing architecture. Given that QoS-based routing is a challenging problem, it may be best to consider the intradomain, interdomain and scalability aspects separately.

REFERENCES

- 1.** R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: An Overview," [RFC 1633](#), July, 1994.
- 2.** L. Zhang, S. E. Deering, D. Estrin, S. Shenker and D. Zappala, "RSVP: A New Resource Reservation Protocol," Technical Report 95-607, ISI, University of Southern California, 1995.
- 3.** J. Moy, "OSPF Version 2," [RFC 1583](#), March, 1994.
- 4.** J. M. Akinpelu, "The Overload Performance of Engineered Networks with Non-Hierarchical Routing," AT&T Technical Journal, Vol. 63, pp. 1261-1281, 1984.
- 5.** G. R. Ash, J. S. Chen, A. E. Frey and B. D. Huang, "RealTime Network Routing in a Dynamic Class-of-Service Network," Proceedings of ITC 13, 1992.
- 6.** D. Mitra, J. Morrison, and K. G. Ramakrishnan, "ATM Network Design and Optimization: A Multirate Loss Network Framework," Proceedings of IEEE INFOCOM '96, 1996.
- 7.** R. Callon et al, "Issues and Approaches for Integrated PNNI" ATM Forum 96-0355, April 1996.
- 8.** B. Rajagopalan, "Efficient Link State Routing," Unpublished Report, available from the author, braja@bellcore.com.
- 9.** H. Ahmadi, J. Chen, and R. Guerin, "Dynamic Routing and Call Control in High-Speed Integrated Networks," Proceedings of ITC-13, pp. 397-403, 1992.
- 10.** Z. Wang and J. Crowcroft, "QoS Routing for Supporting Resource Reservation," available at <http://boom.cs.ucl.ac.uk/staff/zwang/pub.htm>.
- 11.** R. G. Moskowitz, "Why in the World is the Web So Slow?," Network Computing, March, 15, 1996.
- 12.** ATM Forum PNNI subworking group, "Private Network - Network Specification Interface v1.0 (PNNI 1.0)", afpnni-0055.00, March 1996.
- 13.** W. C. Lee, "Topology Aggregation for Hierarchical Routing in ATM Networks," ACM SIGCOMM Computer Communication Review, 1995.
- 14.** Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4),"

Internet Draft, September, 1992.

15. B. Rajagopalan and M. Faiman, "A Responsive Distributed Algorithm for Shortest-Path Routing within Autonomous Systems," Journal of Internetworking Research, pp. 51-69, March, 1991.

16. D. Estrin, T. Li, Y. Rekhter, K. Varadhan, and D. Zappala, "Source Demand Routing: Packet Format and Forwarding Specification (Version 1)," [RFC 1940](#), May, 1996.

[draft-nair-qos-based-routing-00.txt](#)

Page 13

17. I. Castineyra, J. N. Chiappa, and M. Steenstrup, "The Nimrod Routing Architecture," Internet Draft, [draft-ietfnimrod-routing-arch-01.txt](#), February, 1996.

18. S. E. Deering and D. P. Cheriton, "Multicast Routing in Datagram Internetworks and Extended LANs," ACM Transactions on Computer Systems, pp. 85-110, May, 1990.

19. J. Moy, "MOSPF: Analysis and Experience," [RFC 1585](#), March, 1994.

20. A. Ballardie, J. Crowcroft and P. Francis, "Core-Based Trees: A Scalable Multicast Routing Protocol," Proceedings of SIGCOMM '94.

21. S. E. Deering, D. Estrin, D. Farinacci, V. Jacobson, C-G. Liu, and L. Wei, "An Architecture for Wide-Area Multicast Routing," Technical Report, 94-565, ISI, University of Southern California, 1994.

22. B. M. Waxman, "Routing of Multipoint Connections," IEEE JSAC, pp. 1617-1622, December, 1988.

23. X. Jiang, "Path Finding Algorithms for Broadband Multicast," in High Speed Networking, III, Elsevier Science Publishers, 1991.

24. C-H. Chow, "On Multicast Path Finding Algorithms," Proceedings of the IEEE INFOCOM '91, pp. 1274-1283, 1991.

25. S. Sibal and A. Desimone, "Controlling Alternate Routing in General-Mesh Packet Flow Networks," Proceedings of ACM SIGCOMM '95, 1995.

26. D. Mitra and J. B. Seery, "Comparative Evaluations of Randomized and Dynamic Routing Strategies for CircuitSwitched Networks," IEEE Trans. on Communications, pp. 102-116, January, 1991.

27. S. Bahk and M. El Zarki, "Dynamic Multi-Path Routing and How it Compares with Other Dynamic Routing Algorithms for High Speed Wide Area Networks," Proceedings of SIGCOMM '92, pp. 53-64, 1992.

28. R. Gawlick, C. R. Kalmanek, and K. G. Ramakrishnan, "On-Line Routing of Permanent Virtual Circuits," Computer Communications, March, 1996.

AUTHORS' ADDRESSES

Bala Rajagopalan
Bellcore
445 South Street, Rm 1A-257B
Morristown, NJ 07960
U.S.A

Ph: +1-201-829-4261
Email: braja@bellcore.com

Raj Nair
Ascom Nexen
289 Great Rd.
Acton, MA 01720
U.S.A

Ph: +1-508-266-4536
Email: nair@nexen.com