Quality-of-Service Routing without Global Information Exchange

Srihari Nelakuditi[†], Rose P. Tsang[‡] and Zhi-Li Zhang[†]

[†] Dept. of Computer Science & Engineering
University of Minnesota
Minneapolis, MN55455
{srihari,zhzhang}@cs.umn.edu

[‡] Sandia National Laboratories
PO Box 969, Mail Stop 9011
Livermore, California 94550
rtsang@ca.sandia.gov

Abstract— **In this paper we propose a novel** ***localized*** **QoS routing approach, under which** ***no global QoS state information exchange among network nodes is needed.*** **This approach has many advantages over the conventional** ***global*** **QoS routing approach based on information exchange among routers. We develop two localized QoS routing schemes and demonstrate through simulations that the proposed localized QoS routing approach is indeed a viable alternative to the global QoS routing approach.**

I. INTRODUCTION

In Quality-of-Service (QoS) based routing, paths for flows are selected based upon knowledge of the resource availability (referred to as *QoS state*) at network nodes (routers or switches) and QoS requirement of the flows. The majority of previous work (see, e.g., [7], [2], [3], [4], [5], [1], [6]) have viewed QoS routing as an extension of the current Internet routing paradigm where routers exchange their QoS states (e.g. *available bandwidth* of network links) using an OSPF-like protocol (we refer to this routing approach as the *global* QoS routing approach). These QoS states form a *current* view of the network QoS state based on which a source node dynamically determines the "best" feasible path for a flow originating from it to a destination node. A well-studied algorithm for selecting the "best" feasible path is the *widest shortest path* (*wsp*) algorithm which, for a given flow, selects the shortest feasible path with the largest amount of available bandwidth.

Because network resource availability changes with each flow arrival and departure, maintaining *accurate* network QoS state requires *frequent* information exchanges among the network nodes. As the flow activity increases, the cost of propagating the correspondingly large number of QoS state updates throughout the network becomes prohibitive. This precludes the possibility of always providing each node with an accurate view of the current network QoS state. This raises an important issue inherent in the global QoS routing approach: *the network QoS state information acquired at a source node can quickly become out-of-date when the QoS state update interval is large relative to the flow dynamics.* Under these circumstances, exchanging QoS state information among network nodes is superfluous. Second, path selection based on a *deterministic* algorithm such as Dijkstra's shortest path algorithm where *stale QoS state information is treated as accurate* does not seem to be judicious.

In order to circumvent this problem, we propose a novel *localized* QoS routing approach. Under this approach, *no global QoS state information exchange among network nodes is needed.* Instead, source nodes infer the network QoS state based on call blocking statistics collected locally. With the assistance of call admission control (CAC) module, these call blocking statistics can be gathered either at the link-level or at the route-level (see Section III). Due to the nature of collected statistics, the path selection algorithm used at source routers is *non-deterministic.* We develop two localized QoS routing schemes using *probabilistic* path selection algorithms: the *link-level random path* (lrp) algorithm uses link-level statistics, and the *route-level random path* (rrp) algorithm uses route-level statistics. We show through extensive simulations that the two localized QoS routing schemes achieve performance comparable to (if not better than) global QoS routing schemes such as *wsp*, when the QoS state update interval for the latter is large relative to the flow dynamics. These results show that the proposed localized QoS routing approach is a viable alternative to the conventional QoS routing approach based on global information exchange.

The proposed localized QoS routing approach has several advantages. First of all, without the need for global information exchange, the communication overhead involved is minimal. Second, core routers (i.e., non-source routers) do not need to keep and update any QoS state database necessary for QoS routing, thereby reducing the processing and memory overhead at core routers. Last but not the least, the localized QoS routing approach does not require any modification or extension to existing routing protocols such as OSPF. Only source routers need to add a QoS routing enhancement to the existing routing module. This makes localized QoS routing schemes readily deployable with relatively low cost.

The rest of the paper is organized as follows. In Section II, we compare *wsp* with a simple random scheme and demonstrate the limitations of the global QoS routing approach in the presence of inaccurate information. In Section III we present the two localized QoS routing schemes and their performance results. Section IV concludes the paper.

II. LIMITATIONS OF GLOBAL QOS ROUTING APPROACH

In this section we use the widest shortest path (*wsp*) algorithm as an example to illustrate the inherent problem of the global routing approach *when QoS routing information becomes inaccurate.* We show that as the degree of inaccuracy in network nodes' QoS state information increases, the performance of *wsp* degrades quickly to that of a naive random routing scheme, referred to as *uniform random path* (*urp*). The *urp* scheme randomly selects a path from a set of pre-computed paths *without* using any QoS state information. Before we proceed, we first describe the simulation setting used in our study.

Figure 1 shows the (*isp*) topology of an ISP backbone network used in our simulations (also used in [1], [4]). For simplicity, all the links are assumed to be bidirectional and of the same capacity, with C units of bandwidth in each direction. Flows arriving into the network are assumed to require one unit of bandwidth. Hence each link can accommodate at most C flows simultane-

0-7803-5671-3/99/$10.00 © 1999 IEEE

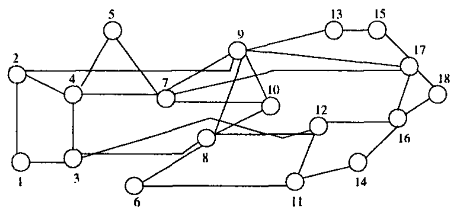

Fig. 1. The *isp* topology used in the simulations

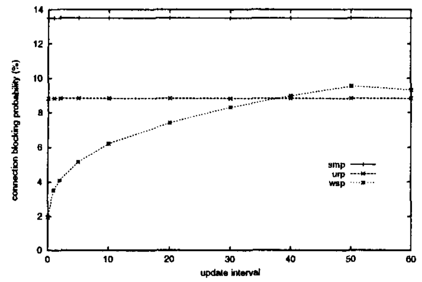

Fig. 2. Illustration of the impact of inaccuracy on wsp (load = 0.65)

ously. The flow dynamics of the network is modeled as follows (similar to the model used in [6]). Flow connection requests arrive at a source node according to a Poisson process with rate λ. The destination node of a flow is chosen randomly from the set of all nodes except the source node. The holding time of a flow is exponentially distributed with mean $1/\mu$. Following [6], the offered network load is given by $\rho = \lambda N\bar{h}/\mu LC$, where N is the number of source nodes, L the number of links, and $\bar{h}$ is the mean number of hops per flow, averaged across all source-destination pairs. The parameters used in simulation are $C = 20$, $\lambda = 1$, $N = 18$, $L = 60$ and $\bar{h} = 2.36$. The average flow holding time $1/\mu$ is set depending upon the desired load.

Figure 2 compares the performance of *wsp* and *urp* as a function of the interval between periodic QoS state updates in *wsp*. In the figure we also plot the performance of a static routing scheme *smp* (static min-hop path), which uses a *fixed* min-hop (i.e. shortest) path for any flow between a pair of source and destination nodes. The performance is measured in terms of the overall connection request blocking probability, which is defined as the ratio of the total number of blocks to the total number of flow arrivals. From the figure we see that *urp* has an overall blocking probability around 9% under the given load $\rho = 0.65$, whereas *smp* has an overall blocking probability close to 14%. In comparison, the overall blocking probability of *wsp* is 2% with the accurate network QoS state (i.e., update interval of 0). However, as the update interval of *wsp* increases, the performance of *wsp* degrades rapidly and approaches that of *urp*. When the update interval goes beyond 40, *wsp* performs worse than *urp*. To put the performance results in perspective, if on an average 10 flows arrive at a source node every second, then update interval of 40 means that global information exchange occurs every 4 seconds. Hence in this case when the update interval is larger than 4 seconds, the performance of *wsp* is worse than that of *urp*, which requires no global information exchange at all. The rapid degradation in the performance of *wsp* in the presence of inaccurate QoS routing information reveals two limitations of the OSPF-like global QoS routing approach: ***the futility of exchanging QoS routing information when the QoS update interval is large relative to the flow dynamics*** and ***the inadequacy of deterministic path selection algorithms in presence of inaccurate QoS routing information.***

III. LOCALIZED QoS ROUTING SCHEMES

In this section, we examine two localized probabilistic QoS routing schemes: the *link-level random path* (lrp) algorithm which uses localized link-level statistical metrics collected by source nodes, and the *route-level random path* (rrp) algorithm which uses route-level statistical metrics.

A. Link-level statistics based QoS routing

In *lrp*, each source node collects link-level statistics locally and performs path selection based on this information. The link-level statistics collected by source nodes are *link blocking probabilities* (lbp's). The *lbp* of a link at a source node is the ratio of the number of connection requests originating at that node which are blocked at the link to the total number of connections attempted through the link from that node. The *lbp* is used as an indicator of how busy the link is. In order to collect this information, we assume that whenever a connection setup request fails at a link, the identity of the link is returned back to the source node, in a failure notification message. The call admission control module at the source node then informs the routing module of the connection setup failure and the identity of the link where the failure occurred. Every source node maintains, for each link l, the number of connection setup requests that are attempted through that link, n_l, and the number of connection setup failures at that link, f_l. The blocking probability of link l, lbp_l, is thus f_l/n_l. We denote the success probability of link l by lsp_l where $lsp_l = 1 - lbp_l$.

Based on *lsp* information, the *lrp* scheme uses a probabilistic path selection algorithm which attempts to pick the lesser congested paths with higher probability. The algorithm works as follows. When selecting a path between a source node s and a destination node d, all paths from s to d are considered in the order of their lengths. Paths of equal length are ordered according to their *psp*'s (path success probabilities), where *psp* of a path is defined as the product of the *lsp*'s of all its constituent links. The shortest path with the largest *psp* is considered first. A path is selected probabilistically through a series of link-level decisions in the following manner: each link l in a given path is selected with probability lsp_l. A path is selected only if all of its constituent links are selected. If any link is rejected, not only is the path under consideration rejected, any other path sharing this link is also rejected. Thus, *lrp* attempts to minimize the resource usage by favoring shorter paths while probabilistically distributing the load across the network in accordance with availability of links.

B. Route-level statistics based QoS routing

In *rrp*, each source node performs path selection based on locally collected route-level statistics. Each source node collects *path blocking probabilities* (pbp's) of paths from that node to the destination nodes. For a path r from the source node to a destination, suppose n_r is the number of connection set-up requests

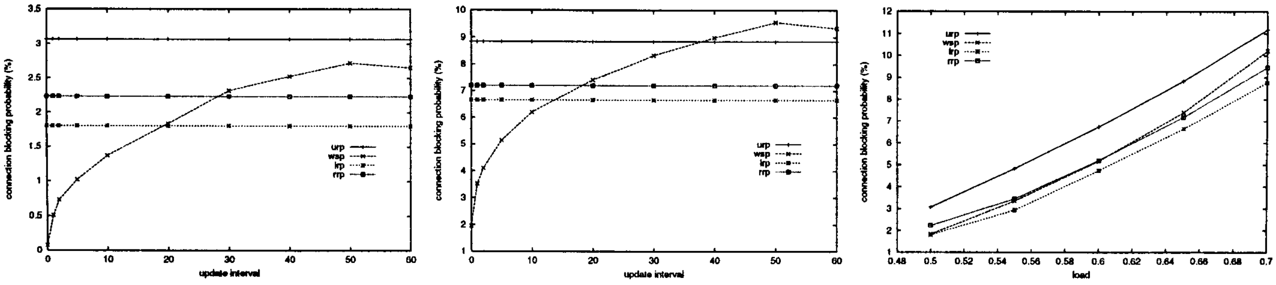

(a) Impact of update interval, $\rho = 0.50$

(b) Impact of update interval, $\rho = 0.65$

(c) Impact of load, update interval of 20

Fig. 3. Comparison of the QoS routing schemes

routed along the path, and b_r is the number of connection set-up requests which are blocked at any link along the path. Then blocking probability of the path, $pbp_r = b_r/n_r$.

Path selection in *rrp* is done in two steps: first a path length is chosen, then a path of the chosen length is selected. Path length is selected based on its *hop success probability*. Let $\mathcal{P}_h$ be the set of paths of length h between a given pair of source and destination nodes. We define *hop success probability* of hop count h, hsp_h, as sum of the path success probabilities of all paths in $\mathcal{P}_h$, i.e., $hsp_h = (\sum_{r \in \mathcal{P}_h}(n_r - b_r))/(\sum_{r \in \mathcal{P}_h} n_r)$. Path lengths are considered in increasing order and so *min-hop* paths are considered first. Let h_0 be the length of the min-hop paths. With probability hsp_{h_0}, we will select a path from the set of min-hop paths. Otherwise, we will decide whether to consider *alternate* paths, i.e., longer paths (than the min-hop paths). This decision is controlled by a configurable parameter α (in our simulations, $\alpha = \frac{1}{2}$). With probability α we will proceed to consider alternate paths and with probability $1 - \alpha$ we will revert back to select a min-hop path. The same procedure is followed for paths of length $h_0 + 1$ and so on. Using such a procedure, *rrp* selects alternate paths with higher probability when min-hop paths are congested but the overall load is low and sticks to min-hop paths when the overall load is high.

Once the path length is decided, a path of given length is chosen with probability *proportional to the inverse of its blocking probability* as follows. Consider the set $\mathcal{P}_h$ of paths of the chosen length h. Let r be a path in $\mathcal{P}_h$. The *path attempt probability* of r, pap_r, is defined as $pap_r = \frac{\frac{1}{pbp_r}}{\sum_{r' \in \mathcal{P}_h} \frac{1}{pbp_{r'}}}$. The path r in $\mathcal{P}_h$ is selected with probability pap_r.

C. Performance evaluation and analysis

Here we compare the performance of the above two schemes with *wsp*. Figure 3 shows the blocking performance comparison of these schemes using the *isp* topology. The overall blocking probability is plotted as a function of the update interval under two load conditions $\rho = 0.50$ (Figure 3(a)) and $\rho = 0.65$ (Figure 3(b)). Figure 3(c) shows the performance of these schemes as a function of the offered network load with update interval set to 20. Note that the update interval is used only in *wsp*.

As expected, the performance of *wsp* degrades quickly as the update interval increases. In addition, at higher loads, the performance of *wsp* degrades faster. We also see that the performance of *lrp* is comparable to that of *wsp* even for small update intervals, particularly at higher loads. Similarly, *rrp* also performs quite well although its performance is worse than that of *lrp*. Both localized probabilistic routing schemes outperform *wsp* when the update interval for *wsp* is large (30 or beyond under $\rho = 0.5$ and 20 or beyond under $\rho = 0.65$). They consistently outperform the simple random scheme *urp* under all load conditions. From the above simulation results, we see that the localized QoS routing schemes yield robust performance without requiring global exchange of the network QoS state.

IV. CONCLUSIONS

We proposed a novel *localized* QoS routing approach that doesn't require global QoS state information exchange among network nodes. This approach has many advantages over the conventional global QoS routing approach based on information exchange among routers. We developed two localized QoS routing schemes, one based on link-level statistics, another based on route-level statistics. We demonstrated through simulations that the two localized routing schemes achieve performance comparable to (if not better than) the global QoS routing scheme *wsp*, when the QoS state update interval for the latter is large relative to the flow dynamics. These results show that the proposed localized QoS routing approach is a viable alternative to the conventional QoS routing approach based on global information exchange among routers.

REFERENCES

[1] G. Apostolopoulos, R. Guerin, S. Kamat, S. Tripathi, "Quality of Service Based Routing: A Performance Perspective", ACM SIGCOMM 1998.

[2] R. Guerin, S. Kamat, A. Orda, T. Przygienda, D. Williams, "QoS Routing Mechanisms and OSPF Extensions", *Work in Progress*, Internet Draft, March 1997.

[3] R. Guerin, A. Orda, "QoS-Based Routing in Networks with Inaccurate Information: Theory and Algorithms", IEEE Infocom 1997.

[4] Q. Ma, P. Steenkiste, "On Path Selection for Traffic with Bandwidth Guarantees", IEEE ICNP 1997.

[5] Z. Zhang, C. Sanchez, B. Salkewicz, E. Crawley, "Quality of Service Extensions to OSPF", *Work in Progress*, Internet Draft, September 1997.

[6] A. Shaikh, J. Rexford, K. Shin, "Evaluating the Overheads of Source-Directed Quality-of-Service Routing", ICNP 1998.

[7] Z. Wang, J. Crowcroft, "Quality-of-Service Routing for Supporting Multimedia Applications", IEEE JSAC Sept 1996