

Should QoS routing algorithms prefer shortest paths?

Karol Kowalik and Martin Collier

Research Institute for Networks and Communications Engineering (RINCE)

Dublin City University, Dublin 9, Ireland,

Telephone: +353 1 700 5805

Fax: +353 1 700 5508

Email: {kowalikk, collierm}@eeng.dcu.ie

Abstract—Multimedia traffic and real-time e-commerce applications can experience quality degradation in traditional networks such as the Internet. These difficulties can be overcome in networks which feature dynamically set up paths with bandwidth and delay guarantees. The problem of selecting such constrained paths is the task of Quality of Service (QoS) routing. This paper considers link-state routing, and the choice of cost metric used to implement QoS routing.

There are two schools of thought regarding the choice of link cost. It is commonly assumed that QoS routing algorithms should limit hop count so as to conserve resources for future connections. Others advocate load balancing mechanisms so as to increase overall network utilisation. This paper investigates which of these approaches gives the better performance. We show that there is no one general answer to this question. We also point out the dangers of drawing general conclusions about routing algorithm performance based on the study of only a limited set of network topologies.

I. INTRODUCTION

Retrofitting the Internet with QoS capabilities is a challenging task. A major concern is that in the current Internet data packets belonging to the same flow may follow different paths to the destination. Standard best-effort service does not provide any guarantees for the traffic flows and so is not suitable for use by multimedia or real-time e-commerce applications. To provide QoS guarantees new service models [1], [2] and mechanisms [3] need to be implemented. Also routing protocols must be enhanced or replaced by QoS routing mechanisms to enable the selection of paths with specific QoS requirements.

There is a wide variety of proposed solutions to the problem of selecting a path with specific QoS requirements - a comprehensive survey can be found in [4] and [5]. This paper considers link-state routing, where the source router selects a route based on information about network resources provided by link state advertisements and about QoS requirements carried within set-up requests. Flooding of link state information is used to ensure that all routers process the same topological and state information. Since it is difficult to maintain up-to-date global state information at every router, QoS routing algorithms should be able to tolerate imprecise information.

Link state routing is normally based on the least cost path computation, although the link cost metric can be expressed in many ways. In this paper we consider mechanisms used to find a path with specific bandwidth requirements, where path cost is usually expressed as a function of the link utilisation. Depending on how link cost metric is formulated these QoS

routing algorithms can address different performance issues (conserving resources, increasing overall network utilisation, etc.). Our investigations suggest that researchers in this field take one of two contradictory approaches: either resource conserving (RC) or load distributing (LD). The RC approach is widely recommended [6], [7], [8], [9] and it has even been proposed by the IETF [10] as a way to introduce QoS routing. Its proponents claim that [6]:

... previous comparative studies have demonstrated that algorithms with a strong preference for minimum-hop routes almost always outperform algorithms that do not consider path length ...

The second approach is recommended in [11], [12], [13] because, as claimed in [14]:

... the use of min-hop routing for on-line routing of permanent virtual circuits can lead to inefficiencies despite the fact that it uses the least amount of resources.

In this paper we investigate the merits of these apparently contradictory claims.

II. THE RESOURCE CONSERVING APPROACH (RC)

There is a large community of researchers who recommend RC [6], [7], [8], [9]. Routing algorithms usually conserve resources by minimising the hop number during path selection process, and calculate a path in two steps: first they select the set of shortest paths; secondly, if there is more than one such path, the criterion of available bandwidth is used to choose among the set of selected paths. As an example of the RC approach we have chosen the *widest-shortest path* (WSP) algorithm [15], which selects the path with the highest residual bandwidth from the set of shortest paths (we assume that link pruning is disabled).

The RC approach has the following advantages:

- By the use of shortest paths it conserves resources – because longer paths use extra resources and can block other potential future requests;
- The primary link cost metric (that used in the first step of path computation) is a constant, and so is robust to network state information inaccuracies.
- Traffic fluctuations are less likely to appear over the set of shortest paths than for the set of all available paths which could be offered by other routing mechanisms. Traffic

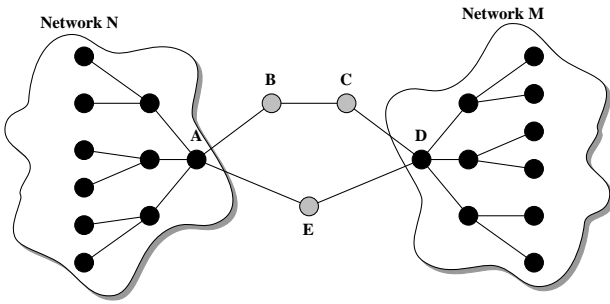


Fig. 1. Example of situation when the use of shortest path is not beneficial

fluctuations occur because links advertised as having a low utilisation will probably be favoured by many incoming requests (they will attract new connections) and in the next update interval this will result in a high load over these links. After advertising a high utilisation few connections will be routed via these links and the whole cycle will repeat.

Despite these characteristics there is one major drawback of this approach: when limiting hop count (and when link pruning is disabled) we can only use links belonging to shortest paths, so that we cannot fully benefit from existing longer and lightly loaded routes. Let us consider the situation shown in Figure 1, where we have two networks N and M interconnected using a short (two hop) path (A-E-D) and a longer path (A-B-C-D). If when establishing connections between network N and M we restrict ourselves only to the use of the shortest path (link pruning being disabled), the longer path remains unused. In such cases RC approach is not beneficial compared even with static routing. We show in Section IV-C that this happens usually for networks with low connectivity.

We can avoid scenarios as described above, simply by pruning links with insufficient resources first and calculating the min-hop path over a reduced topology. However it was observed [16], [6] that when the state information is highly imprecise it is difficult to decide whether or not link has enough resources to accommodate the new request. The pruning process can remove infeasible as well as feasible links, so it degrades the overall performance. In this paper we assume that inaccuracy of state information is present at all times, so pruning is disabled.

For some networks (such that shown in Figure 1) the inefficient usage of resources by the RC approach can be overcome by admitting paths with hop count equal to the minimum plus one [9], [15]. However in this paper we are looking for general conclusions, so such ad-hoc solutions are not considered here.

III. THE LOAD DISTRIBUTING APPROACH (LD)

Considerable attention has also been given to LD algorithms [11], [12], [13]. Such load balancing is often achieved by the use of link cost metrics expressed as a convex function which increases with the link utilisation [11], [12], [13], [17]. As an example of this approach we have chosen the *shortest-distance path* (SDP) algorithm [15], which computes the least

cost path, where link cost is inversely proportional to available bandwidth. The benefit of using such a convex function for the link cost is intuitively clear: it favours paths on lightly loaded links over those on busy links, so that the load is balanced over all links.

This approach to QoS routing has following advantages:

- It avoids congested links;
- It distributes load over the network, so that it can fully exploit longer paths;

It also has a number of drawbacks:

- By using longer paths it consumes extra resources, and this can block future requests. However this is not a crucial issue when the network load is light [15].
- Due to the use of a main cost metric which depends on information about available resources, it requires up-to-date state information. Otherwise considerable performance degradation is observed. In [18] we have shown a way to tackle this problem.

QoS routing employing LD can be beneficial in situations such as shown in Figure 1, but in general it is vulnerable to inaccuracy of network state information, which can drastically degrade its performance.

IV. PERFORMANCE EVALUATION

In this section we compare the performance of the above two approaches to QoS routing. In order to draw general conclusions about their performance, we study their use on a wide variety of network topologies.

A. Network Model

Researchers have evaluated the performance of different routing approaches on simulated networks with various topologies such as: ISP [19], [5], MCI [15], random topologies [8], [6], regular topologies [6], and others [14], [9]. However the topology of the Internet is difficult to characterise because it is constantly changing [20], so the use of a limited set of topologies should not be recommended.

We want to show how routing algorithm performance is affected by the level of network connectivity. Thus we require a network topology whose connectivity can be adjusted in a controlled manner. To achieve this, we decided to use a base network with N nodes and L links and to add links randomly until the required connectivity was achieved. As a base network we use the topology shown in Figure 2, which has $N = 21$ nodes and $L = 20$ links (the minimal level of connectivity corresponding to $L = N - 1$), so the average node degree is equal to $2L/N = 1.9$. We add links randomly to our base network until the number of links is $L = 80$ which gives an average node degree of 7.6.

The network connectivity is an important factor when the LD approach is evaluated, while when the RC approach is used a more meaningful parameter is the number of min-hop paths, so for each connectivity level we also present the average number of min-hop paths between any source and destination pair.

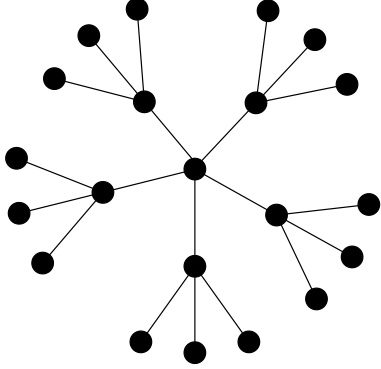


Fig. 2. The base network topology

Results obtained by other researchers usually relate to topologies with a degree of around four, and rarely report results for degrees higher than six and less than tree (e.g., [8], [6]). We study a wider range of connectivities.

We do not consider regular topologies such as the cube [6], because they are infrequently used in real networks. They feature many min-hop paths between source and destination, which has a positive impact when RC approach is used. However such regularities are not present in real networks.

We extend the base network by adding links until rich connectivity is achieved. It would also be possible to remove links from a highly connected network to obtain lower connectivity. The mechanism used to deal with broken connections could have a significant impact on overall performance increasing the variability of results obtained. So we decided that adding links to a base network of low connectivity is the better approach.

Adding one or more links to a network can result in different routing performance, depending on the placement of the newly added links. Therefore this process needs be repeated several times to get a stable and uniform result. In our experiment we create 200 different network topologies with the same connectivity level, and average their performance. All the average results are obtained with a 95% confidence interval. We also assume that links are bidirectional, each with identical capacity C (in our simulations we use DS-3 links with capacity $C = 45Mbps$).

B. Traffic model

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 $B = 3.32Mb/s$. If N^a nodes in the network generate the traffic, the load offered to the network is [6] $\rho = \lambda N^a B h' / \mu LC$, where h' is the average shortest path distance between nodes, calculated over all source-destination pairs (for our simulation: $N^a = 21$, and $h' = 3.04762$ for the base network shown in Figure 2). In our experiment we adjust λ to produce the required offered load and fix the mean connection holding time at 180sec. However

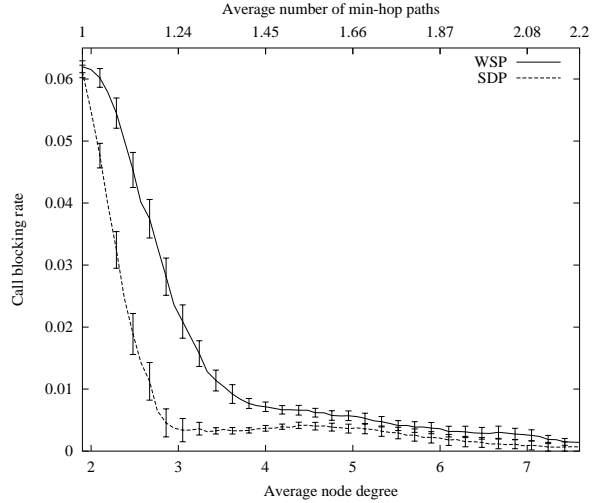


Fig. 3. Call blocking probability of WSP and SDP under increasing node degree, $\rho = 0.3$, $hd = 1sec$

each time a link is added, the values of L and h' change. To keep the network load ρ constant, we recalculate the values of L , h' and λ every time a new link is added.

C. Results

The goal of our experiment is to compare the influence of both network connectivity and inaccuracy of state information on the performance of both RC and LD. We change the level of connectivity as described in Section IV-A, and introduce inaccuracy through the use of *hold-down* timers which set the minimal time interval between concurrent link state updates (we will denote this interval as hd). We assume that each node uses a threshold based update policy [19], [5] with the threshold trigger of 10% controlled by a hold-down timer. We use only one hold-down timer for all links outgoing from the node.

We compare the performance of RC and LD approaches using the *call blocking rate* – defined as:

$$\text{call blocking rate} = \frac{\text{number of rejected requests}}{\text{number of arrived requests}}$$

The first step of our experiment compares both algorithms under a light load, $\rho = 0.3$. As shown in Figure 3 when the load is light and the state of the information is up-to-date ($hd = 1sec$), LD gives better performance. Similar results occur even if the level of inaccuracy in network state information is much higher, for example when $hd = 40sec$ as shown in Figure 4. This confirms the findings presented in [15] and can be explained as follows: using extra resources, in the form of longer paths, has negligible importance when the network load is low, because there are still sufficient uncommitted resources to accommodate future connections. Selecting shortest-path routes cannot fully exploit all network resources, and produces a higher blocking rate. This findings remain true over a wide spectrum of network connectivity (average node degree from 1.9 to 7.6).

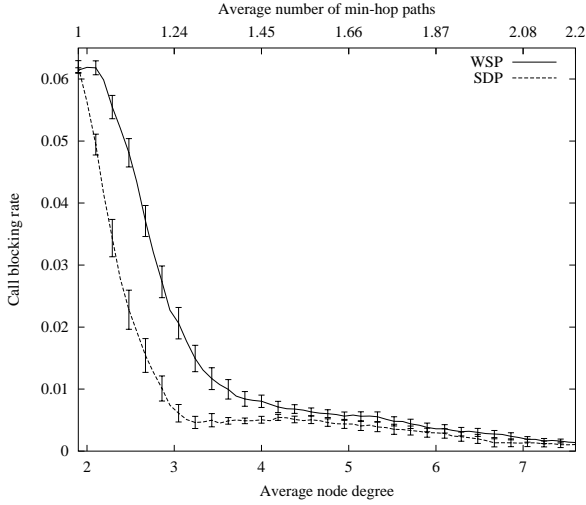


Fig. 4. Call blocking probability of WSP and SDP under increasing node degree, $\rho = 0.3$, $hd = 40sec$

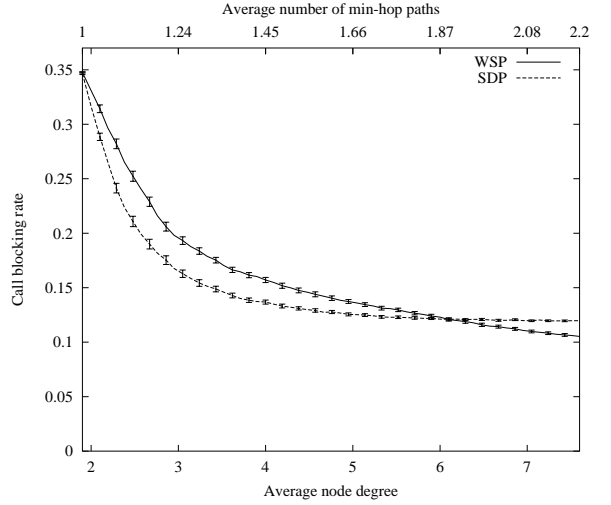


Fig. 5. Call blocking probability of WSP and SDP under increasing node degree, $\rho = 0.8$, $hd = 1sec$

The situation changes if the network load is high. As shown in Figure 5 if accurate state information is provided ($hd = 1sec$), LD makes better use of network resources for networks with node degree less than 6. However if inaccuracy is higher, for example when $hd = 20sec$ or $40sec$ as shown in Figure 6 and Figure 7 respectively, RC approach performs well even for less connected networks (for node degree > 4.5 if $hd = 20sec$ and > 3.5 if $hd = 40sec$). This suggests that load balancing approaches cannot operate well in the presence of very inaccurate state information. The RC approach is more resistant to imprecision of state information, and for a highly connected network this gives a huge advantage over the LD approach.

Nevertheless there are networks with low connectivity, such as the one shown in Figure 1, for which only load balancing approaches can improve on static shortest path routing.

These results show that network topology, traffic and accuracy of the state information play an important role in deciding which routing strategy performs better. Hence none of these three factors can be neglected in an evaluation phase of routing mechanisms.

V. CONCLUSIONS

The respective merits of RC and LD have never been comprehensively assessed in the literature, with many researchers advocating one or the other on the basis of results obtained on specific network topologies. We have investigated both classes of routing algorithm over a wide range of network connectivities, and have made the following observations:

- For networks with low connectivity the LD approach outperforms the approach limiting hop count, because it uses scarce resources in a more efficient way.
- RC algorithms typically perform better for highly connected networks - when there are few min-hop paths between any source and destination.

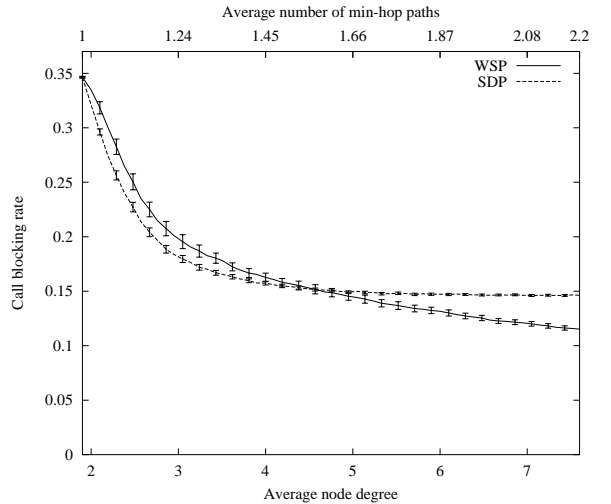


Fig. 6. Call blocking probability of WSP and SDP under increasing node degree, $\rho = 0.8$, $hd = 20sec$

- For lightly loaded networks the LD approach seems to be a much better solution than RC.
- For highly loaded networks the RC approach offers better performance, especially in the presence of very imprecise state information.

Moreover these observations suggest that in order to allow QoS routing to perform well over a wide range of network topologies it cannot be restricted only to the set of shortest paths. This suggests the use of both RC and LD mechanisms, to result in an efficient usage of the resources and low vulnerability to imprecise network state information.

Our findings show the dangers of drawing conclusions about the performance of a routing algorithm on the basis of results obtained using a limited set of network topologies. Because of the wide diversity of network topologies constituting the

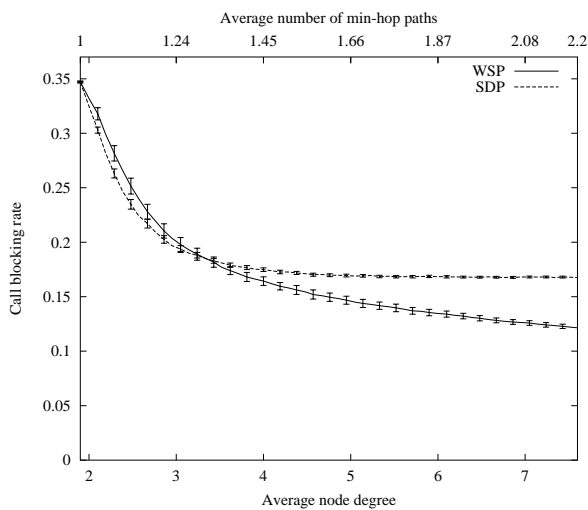


Fig. 7. Call blocking probability of WSP and SDP under increasing node degree, $\rho = 0.8$, $hd = 40sec$

current Internet, routing algorithms should be tested over a wide spectrum of connectivity. Only such an approach allows general conclusions to be drawn about the performance of the algorithm.

REFERENCES

- [1] B. Braden, R. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: An Overview," *RFC 1633*, July 1994.
- [2] D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Services," *RFC 2475*, December 1998.
- [3] B. Braden, R. Zhang, L. Berson, S. Herzog, and S. Jamin, "Resource Reservation Protocol (RSVP) - Version 1 Functional Specification," *RFC 2205*, September 1997.
- [4] S. Chen and K. Nahrstedt, "An Overview of Quality of Service Routing for Next-Generation High-Speed Networks: Problems and Solutions," *IEEE Network*, pp. 64–79, Nov./Dec 1998.
- [5] X. Yuan and W. Zheng, "A Comparative Study of Quality of Service Routing Schemes That Tolerate Imprecise State Information," Florida State University Computer Science Department, Technical Report. [Online]. Available: <http://websrv.cs.fsu.edu/research/reports/TR-010704.pdf>
- [6] A. Shaikh, J. Rexford, and K. G. Shin, "Evaluating the impact of stale link state on quality-of-service routing," *IEEE/ACM Transactions on Networking (TON)*, vol. 9, no. 2, pp. 162–176, 2001.
- [7] Q. Ma and P. Steenkiste, "Quality-of-service routing for traffic with performance guarantees," in *Proc. IFIP International Workshop on Quality of Service*, Columbia University, New York, May 1997, pp. 115–126.
- [8] C. Pornavalai, G. Chakraborty, and N. Shiratori, "QoS Based Routing Algorithm in Integrated Services Packet Networks," in *Proceedings of*

International Conference on Network Protocols, Atlanta, Georgia, October 1997, pp. 167–175.

- [9] I. Matta and A. Shanka, "Dynamic Routing of Real-Time Virtual Circuits," in *Proceedings of IEEE International Conference on Network Protocols*, 1996, pp. 132–139.
- [10] G. Apostolopoulos, D. Williams, S. Kamat, R. Guerin, A. Orda, and T. Przygienda, "QoS Routing Mechanisms and OSPF Extensions," *RFC 2676*, August 1999.
- [11] B. Awerbuch, Y. Azar, S. Plotkin, and O. Waarts, "Throughput-Competitive On-line Routing," *34th Annual Symposium on Foundations of Computer Science*, November 1993.
- [12] R. Gawlick, A. Kamath, S. Plotkin, and K. Ramakrishnan, "Routing and Admission Control in General Topology Networks," *Technical Report STAN-CS-TR-95-1548*, 1995.
- [13] A. Kamath, O. Palmon, and S. A. Plotkin, "Routing and Admission Control in General Topology Networks with Poisson Arrivals," *SODA: ACM-SIAM Symposium on Discrete Algorithms*, 1996.
- [14] R. Gawlick, A. Kamath, and K. Ramakrishnan, "Online routing for virtual private networks," *Computer Communications*, vol. 19, pp. 235–244, March 1996.
- [15] Q. Ma and P. Steenkiste, "On path selection for traffic with bandwidth guarantees," in *Proceedings of IEEE International Conference on Network Protocols*, October 1997.
- [16] G. Apostolopoulos, R. Guerin, S. Kamat, and S. K. Tripathi, "QoS Routing: A Performance Perspective," in *Proceedings of SIGCOMM*, Vancouver, Canada, September 1998.
- [17] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," in *Proceedings of IEEE Infocom 2000*, March 2000, pp. 519–528.
- [18] K. Kowalik and M. Collier, "ALCFRA - A Robust Routing Algorithm Which Can Tolerate Imprecise Network State Information," *15th ITC Specialist Seminar*, July 2002.
- [19] G. Apostolopoulos, 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.
- [20] S. Floyd and V. Paxson, "Difficulties in simulating the Internet," *IEEE/ACM Transactions on Networking*, vol. 9, no. 4, pp. 392–403, August 2001.