Impact of "Trunk Reservation" on Elastic Flow Routing

Sara Oueslati-Boulahia and James W. Roberts

France Telecom-CNET
38-40 rue du Général Leclerc, 92794 Issy-les-Moulineaux Cedex 9, France
{sara.boulahia,james.roberts}@cnet.francetelecom.fr

Abstract. This paper focuses on routing strategies for elastic flows within a flow-based routing architecture. In contrast to other elastic flow routing studies, we assume that elastic flows have a minimum guaranteed rate and are, therefore, subject to admission control. Our ultimate goal is to devise an adaptive algorithm that maximizes the throughput of elastic flows at light load and that preserves network efficiency under overload conditions. To this end, we examine in this paper the impact of trunk reservation on elastic flow quality of service. Trunk reservation is a technique commonly used in circuit-switched networks to prevent performance degradation under overload conditions. Two algorithms integrating a form of trunk reservation are proposed and their performance compared by means of simulation to the performance achieved by an alternative algorithm that we proposed and evaluated in an earlier study [1]. Our results highlight interesting features of trunk reservation when used for elastic flow routing.

1 Introduction

Several studies in the past have considered the problem of routing flows with bandwidth or delay guarantees [2–4]. Recently, there has been a growing interest in addressing issues related to the routing of best effort flows in multiservice networks [1, 5–9]. This paper considers a microflow-based routing architecture for Internet-like networks. More precisely, we assume routing is performed for individual transactions, each transaction being represented as a continuous flow of packets. In practice such a flow might be identified by the five tuple (*source address*, *destination address*, *source port number*, *destination port number*, *protocol identifier*) together with the fact that successive packets are not separated by more than a certain time interval. We do not consider here the feasibility of routing objects defined at such a fine granularity which is clearly an important issue. We seek rather to evaluate the potential advantage of flow sensitive routing algorithms with respect to quality of service and network traffic handling efficiency.

The service model we consider in our routing architecture was proposed in [10] and distinguishes two classes of flow, elastic flows and stream flows. Elastic flows result from the transfer of a digital object (text, image, video) from one place

G. Pujolle et al. (Eds.): NETWORKING 2000, LNCS 1815, pp. 823-834, 2000

to another. They are characterized by a volume (in bytes) and their transfer rate adjusts to the amount of available bandwidth as a result of being controlled by a mechanism such as TCP. Their transfer time depends on the throughput they achieve which thus constitutes their main QoS parameter. Note also that the faster elastic flows complete, the sooner resources are released for ongoing and future flows. Consequently, enhancing elastic flow throughput positively impacts both elastic flow QoS and network resource utilization. Stream flows, on the other hand, mainly result from audio and video applications. Unlike elastic flows, stream flows have an intrinsic rate (possibly variable) and an intrinsic duration and generally correspond to delay sensitive applications.

We assume here that elastic flows are guaranteed a minimum rate by means of admission control. In other words, new flows will not be admitted if the rate of ongoing elastic flows would consequently be reduced below a certain limit. This additional feature impacts the design of elastic flow routing algorithms. The guarantee of a minimum rate to elastic traffic is largely motivated in [5, 6]. The detailed admission control procedure envisaged for a multiservice network integrating elastic and stream traffic, as well as the underlying resource sharing model, are described in [1].

The present study focuses on routing strategies for elastic flows. Our objective is to devise a routing algorithm that enhances as much as possible the throughput of elastic flows under light traffic conditions, while efficiently handling traffic at overload. Initial results obtained in [7] suggest that, in order to maximize the throughput of best effort flows, least loaded paths should be preferred to minimum-hop paths under light load conditions, and that shortest paths should, on the contrary, be privileged under heavy load conditions. Ideally, the path selection algorithm does this automatically without requiring a manually specified threshold that distinguishes between heavy load and light load. In this perspective, we proposed and evaluated in earlier work a novel routing algorithm, called *Maximum Utility Path*, which proved to be particularly robust with respect to topology configurations and traffic conditions [1, 11].

In this paper, we investigate for the first time the benefit for elastic flow routing of the technique widely used in circuit-switched networks to prevent performance degradation under overload conditions known as *trunk reservation*. More precisely, we evaluate the performance of two algorithms that integrate the trunk reservation technique. The latter algorithms derive from two well known algorithms, *Widest-Shortest Path* and *Shortest-Widest Path*. The performance of the proposed algorithms, called *Widest-Shortest Path with Trunk Reservation* and *Shortest-Widest Path with Trunk Reservation*, are compared to that obtained with the *Maximum Utility Path* algorithm, under various load conditions and on different network topologies. In all our simulation scenarios, only elastic traffic is offered to the network.

The next section discusses the benefit of trunk reservation for enhancing elastic flow QoS under heavy traffic conditions. In Section 3, we describe our routing framework: we present the link metrics we consider, discuss admission control issues and describe the routing algorithms to be evaluated. The following section presents our simulation model and in Section 5 simulation results are analyzed. We conclude the paper by a summary of our main simulation results.

2 Trunk Reservation for Elastic Flows

Trunk reservation is an easily implemented control mechanism that is used in circuit-switched networks to protect the network against instability due to the excessive use of alternative paths during overload conditions. An incoming call that is rejected by its direct path is admitted on an alternative path, with two or more hops, only if the number of free circuits on all the links of the path is more than a certain threshold. The implementation of trunk reservation in telephone networks ensures that direct calls are given priority over less efficient overflow calls. One potential concern with the trunk reservation scheme is how to choose the threshold value. Instead of an absolute threshold we use the notion of *Trunk Reservation Level* (TRL) which represents the proportion of capacity that is reserved to priority flows. Several studies [12–14] have attempted to identify optimal values for the TRL which typically represents a small proportion of accessible trunks.

In this paper, we focus on networks carrying elastic flows with a guaranteed minimum rate. In this context, by trunk reservation we mean the reservation of a certain proportion of bandwidth for minimum-hop path routed traffic. The amount of bandwidth that should be reserved is one object of our investigation. Note that in the context of elastic flow routing, the throughput of flows constitutes an additional performance metric to be taken into consideration. In the remainder of this paper *directly-routed traffic* refers to flows routed on their minimum-hop path(s) (not neccessarily direct paths) and *overflow traffic* refers to flows rejected by their minimum-hop path(s).

At this point we investigate the benefit of trunk reservation for elastic flows at the level of a single link. Consider a router outgoing link of capacity C simultaneously used by N elastic flows belonging to two different classes : directly-routed traffic (class 1) and overflow traffic (class 2). The difference resides in the access priority. Flows of the second class have a more restricted access than those of the first class. The link is modeled by a Processor Sharing queue with two thresholds S_1 and S_2 with $S_1 \geq S_2$. This implies that when N flows are in progress, each receives a throughput equal to $\frac{C}{N}$, idealizing the fair sharing objectives of TCP congestion avoidance algorithms. A flow of class i is admitted only if the number of ongoing flows is smaller that S_i, $i = 1, 2$. We take $C = 100$ and set the minimum rate for all flows to 1% of C, so that $S_1 = 100$. The relation between S_1, S_2 and the TRL value is given by: $S_2 = S_1 \times (1 - TRL)$.

We further assume a Poisson arrival process of intensity λ_i for flows in classes i, $i = 1, 2$. The size of flows in both classes has a general distribution with mean $\frac{C}{\mu}$. Let B_i denote the blocking probability for class i. $B_1 = P[N = S_1]$ and $B_2 = P[N \geq S_2]$. Note that the mean sojourn time is the same for both classes, because, once in the system, all flows get equal link bandwidth shares. B_1 and B_2 are given by the following expressions:

$$B_1 = \rho^{S_2} * \rho_1^{S_1 - S_2} * P[N = 0] \quad (1)$$

$$B_2 = \frac{\rho^{S_2} * (1 - \rho_1^{S_1 - S_2 + 1}) * P[N = 0]}{1 - \rho_1} \quad (2)$$

with

$$P[N = 0] = \left(1 + \frac{\rho * (1 - \rho^{S_2})}{1 - \rho} + \frac{\rho_1 * \rho^{S_2} * (1 - \rho_1^{S_1 - S_2})}{1 - \rho_1}\right)^{-1}, \ \rho_1 = \frac{\lambda_1}{\mu}, \ \rho = \frac{\lambda_1 + \lambda_2}{\mu}$$

For offered loads less than 1, both B_1 and B_2 are very small for any reasonable choice of S_2 (e.g. $S_2 \geq 40$). Differentiation is significant mainly in overload conditions ($\rho > 1$). Figures 1 and 2 plot the blocking probability of each class of flow for $\rho = 1.2$ and $\rho = 2.4$, respectively. Flows in both classes have the same arrival intensity ($\lambda_1 = \lambda_2$).

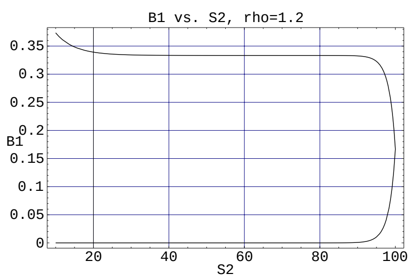

Fig. 1. $\rho = 0.6 * 2 = 1.2$

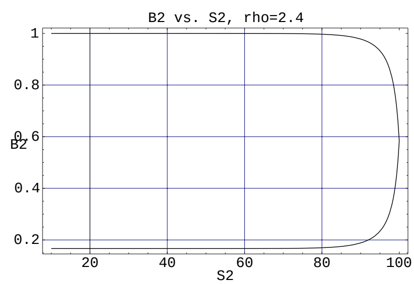

Fig. 2. $\rho = 1.2 * 2 = 2.4$

We observe that offering a limited access to overflow traffic on a link effectively protects flows for which this link composes their minimum-hop path. For a load of 1.2, only overflow traffic is blocked and when the load is high enough to cause blocking to class 1 flows, virtually all class 2 flows are blocked. We also note that this protection is effective for a wide range of threshold values S_2. Blocking probabilities of both classes converge when S_2 approaches the value of S_1.

The use of trunk reservation on any single link of a network has ramifications which may extend to the rest the network. Our objective in the remainder of the paper is to examine the impact of trunk reservation on elastic flow performance at the level of a network. To this end, we consider two algorithms incorporating trunk reservation and compare their performance to that provided by the algorithms from which they derive.

3 Routing Framework

In this section, we describe the link metrics we consider, specify the condition for admitting a new flow and present the routing algorithms to be evaluated.

3.1 Link Metrics and Admission Conditions

In our routing architecture, elastic flows are assured a common minimum rate. The imposition of this minimum rate is intended to avoid the negative impact of very low rates occurring notably in situations of overload and should not be viewed as a customizable service guarantee. Users need not know about it at all.

All the elastic flow routing algorithms described in the following section use the *offered rate* as a link metric. The offered rate represents an equal fair share of the link bandwidth available to elastic flows. In the present study, the link capacity C is entirely dedicated to elastic traffic and the offered rate is given by: $r = C/(N_e + 1)$. Assuming a fair bandwidth sharing flow control mechanism, the fair share rate offered by a given link is tightly correlated to elastic flow QoS. The offered rate on a link actually represents a lower bound on what flows can actually attain due to the rate limitations affecting competing flows on other links. Of course, the new flow will not realize this rate if its bottleneck link is elsewhere. The path offered rate is equal to the smallest offered rate of all the links constituting the path. A path is said to be feasible for an elastic flow if its offered rate is at least equal to the minimum rate.

Note that maintaining a count of ongoing elastic flows assumes that nodes are able to detect a new flow and determine if a flow is still active or has expired. Besides, in a link state approach the metric *current elastic flow count* would need to be advertised using a certain update mechanism. The potential inaccuracy induced by the update mechanism and its impact on routing performance are not considered in this paper. In addition to this dynamic metric, a static route metric representing the number of hops along a route is used to allow resource usage control.

3.2 Description of the Routing Algorithms

We evaluate two novel routing algorithms that integrate a form of trunk reservation: *Widest-Shortest Path with Trunk Reservation* (WS/TR) and *Shortest-Widest Path with Trunk Reservation* (SW/TR). Both algorithms are meant to enhance the performance of the algorithms *Widest-Shortest Path* (WS) and *Shortest-Widest Path* (SW), respectively. An alternative algorithm, called *Maximum Utility Path* (MU) introduced in [1], is also presented with the objective of comparing its performance to that achieved by the two new algorithms.

A number of studies [1, 7, 11] have evaluated the performance of WS and SW for elastic flows. WS selects the feasible path with the largest offered rate among those with the smallest number of hops. This algorithm privileges "cheap" paths to more resource consuming paths with higher offered rates, and performs therefore better at high load than at light load. One concern with this algorithm is its sensitivity to topology configurations which may exacerbate its poor performance under light load conditions by limiting its ability to perform load balancing (due to the absence of several paths which are equivalent in terms of number of hops). SW, on the contrary, selects the path offering the highest rate with the number of hops deciding ties. Results obtained in [1, 7] show that SW yields high per-flow throughput at light load, but its performance rapidly deteriorates as the load increases due to excessive resource consumption.

A possible enhancement to WS and SW is to use "trunk reservation" to prevent the use of long paths when the network is congested. WS/TR (resp. SW/TR) performs like WS (resp. SW) except that non minimum-hop paths are feasible only if a certain proportion of bandwidth remains available on their bottleneck link to "directly-routed" flows.

MU, on the other hand, evaluates the utility of routing a flow on each path and choses the path with maximum utility. The following utility function is defined:

$$U(P, B) = log(r) - B * \Delta h \tag{3}$$

where r represents the offered rate along path P, Δh represents the difference in the number of hops between path P and the shortest feasible path, and B is a constant (set to 1 in this study). The utility of a route increases with r and decreases linearly with the number of extra hops Δh. An interesting property of this algorithm is its overflow condition, i.e., the condition for selecting a route longer than the shortest feasible one. Consider two feasible candidate paths P_0 and P_1 with h_0 and h_1 hops, respectively. Assume P_0 is the shortest feasible route. Let r_0 and r_1 denote the respective bottleneck rate for each route. The algorithm will select P_1 instead of P_0 if: $\frac{r_1}{r_0} > e^{B*(h_1-h_0)}$. From this condition, it appears that the probability of selecting a path longer than the shortest feasible path decreases exponentially with the number of extra hops on that path. Note that by varying B from 0 to ∞, we cover the spectrum between the *Widest Path* algorithm (with $B = 0$), and the *Minimum-Hop Path* algorithm (with $B = \infty$). Another interesting property of this algorithm is that it automatically places more emphasis on resource conservation as load increases due to the strict concavity of the log function.

4 Simulation Model

Our evaluation is based on the use of an event-driven simulation that performs path selection, admission control and bandwidth reservation operations at the flow level. In the simulations, an elastic flow has an instantaneous throughput equal to the bottleneck offered rate of its route. This choice is less efficient than max-min fair bandwidth sharing but preserves its essential properties while reducing simulation complexity.

Figure 3 illustrates the network topologies considered in this study. The Europe and US topologies are derived from actual commercial backbone networks. The topologies considered exhibit different characteristics with regard to size, the degree of connectivity (the US topology is loosely meshed, the Europe topology is rather tightly meshed and the four-node topology is completely meshed) and symmetry.

The traffic model considers elastic traffic only. We assume a Poisson arrival process for elastic flows with random selection of the source-destination pair

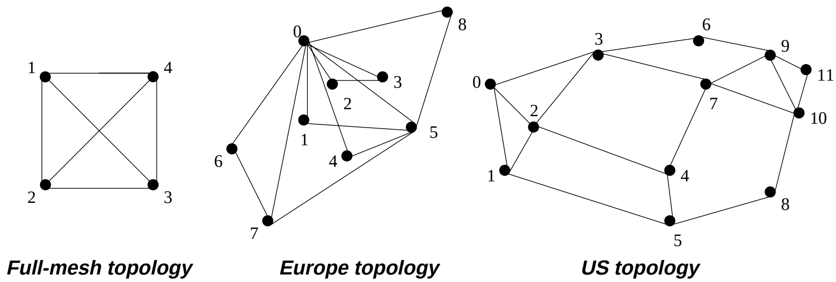

Fig. 3. Network topologies considered

according to equal probabilities (uniform traffic distribution). Elastic flows are characterized by a minimum rate equal to 1 BU, and an access rate which may limit their throughput. We further assume that elastic flow size follows a Pareto distribution with parameter $\alpha = 1.5$. With the Pareto parameters we consider, the resulting average size of flows is 0.66 BU second, and 90% of elastic flows have a size less than or equal to 1 BU second.

Elastic traffic performance is measured jointly by the average per-flow throughput and the blocking probability. We present results below in terms of the *fractional average throughput*: the ratio of the average per-flow throughput to the link capacity (100 BUs).

5 Simulation Results

We first present results when flows have no rate limitation outside the considered networks. We next study the impact of a limited access rate on the performance of the routing algorithms. In the remainder of the paper, alternative paths for a given flow correspond to the subset of candidate paths with at least one hop more than the minimum-hop path(s).

5.1 Comparative Study with no Access Rate Limitation

Figure 4 plots the fractional average throughput as a function of the load, as achieved by WS, MU, SW, WS/TR with a TRL of 0.05 and SW/TR with a TRL of 0.99. Other values of TRL were considered for SW/TR (0.9, 0.5 and 0.1); Figure 6 shows the resulting fractional average throughput for every TRL tested. We also considered two other values of TRL for WS/TR (0.1 and 0.25). We do not show the corresponding throughput results, because hardly any difference was noticed compared to WS/TR(0.05). Table 1 gives the blocking rates produced by all the algorithms.

We first observe that SW/TR(0.99) performs best in terms of blocking and throughput under all load conditions. Note that a TRL of 0.99 in our simulations means that the overflow traffic on a given alternative path is limited to

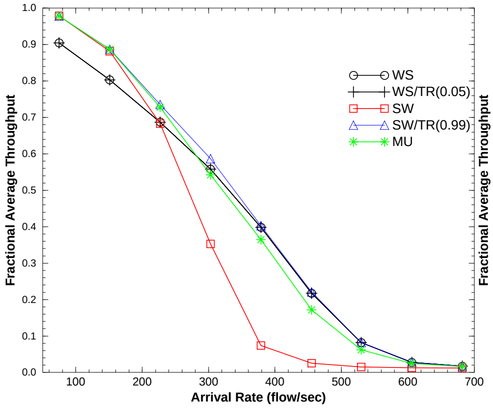

Fig. 4. Throughput comparison with a limited access rate, Full-mesh topology

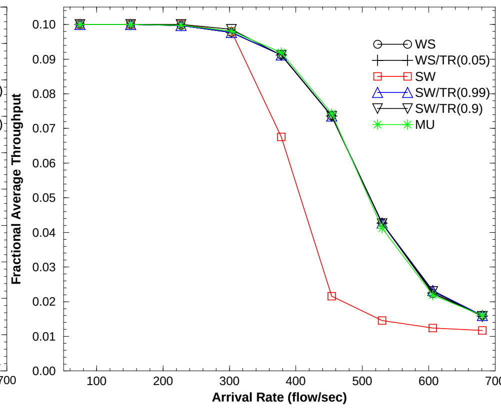

Fig. 5. Throughput comparison with $r^+ = 10BUs$, Full-mesh topology

Table 1. Blocking rates, Full-mesh topology

Arrival rate	MU	WS	WS/TR(.05)	WS/TR(.25)	SW/TR(.99)	SW/TR(.5)	SW
530 flows/s	0 %	0%	0%	0%	0 %	0.2%	9.4 %
606 flows/s	5.8%	6.2 %	2.5 %	2.1 %	2.1 %	4.1 %	19.3 %
681 flows/s	17.0%	16.5 %	9.7 %	8.9 %	8.9 %	10.8 %	25.7 %

a single flow. This appears very restrictive and yet yields good performance. The reason is that, SW/TR(0.99) adopts, as the load increases, virtually the same behavior as the *Minimum-Hop Path* algorithm which, on fully-connected uniformly loaded networks, yields the best possible performance at high load. Figure 6 and Table 1 further show that the performance of SW/TR vary significantly with the TRL value. We also noticed that trunk reservation, even with small TRL values, effectively reduces blocking rates compared to that obtained with SW (e.g. 1.78% for a TRL=0.1 at a rate of 530 flows/s). The throughput, however, is only enhanced for high TRL values (e.g. 0.9 and 0.99). MU, on the other hand, performs as well as SW/TR(0.99) at light load, and moves close to WS as the load increases. As far as WS/TR is concerned, figures in Table 1 show that trunk reservation significantly decreases blocking while having very little effect on throughput (same as WS). At low load, trunk reservation is not activated and WS/TR yields the same performance as WS. At overload, both algorithms cannot offer better rates than the minimum rate, however more flows are admitted with trunk reservation.

Results relative to the Europe topology (Figures 7 and 8 for throughput, Table 2 for blocking) highlight a different effect of trunk reservation compared to that obtained on the full-mesh topology. Regarding the performance of SW/TR, we observe that high TRL values (e.g. 0.99 and 0.9) result in a much higher throughput than the other algorithms, but at the cost of higher blocking rates.

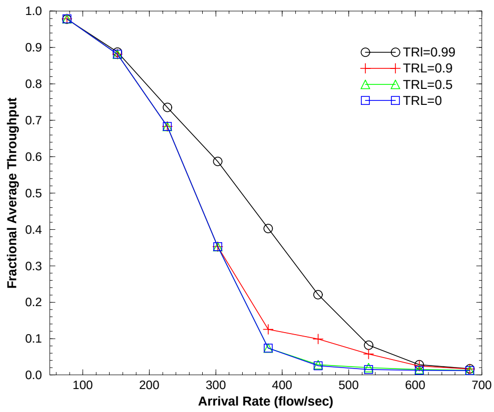

Fig. 6. SW/TR, Full-mesh topology

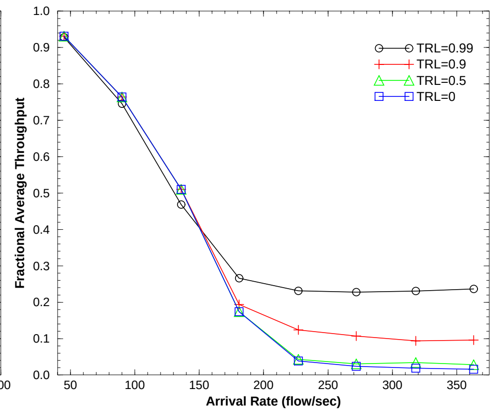

Fig. 7. SW/TR, Europe topology

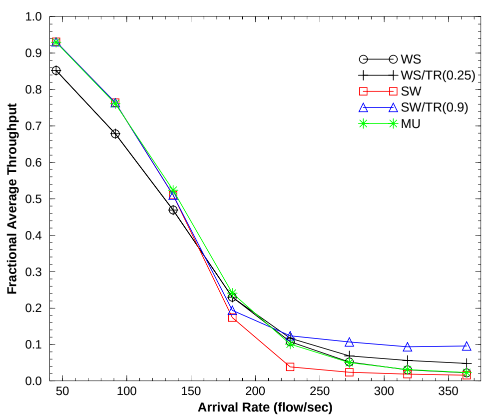

Fig. 8. Throughput comparison, Europe topology

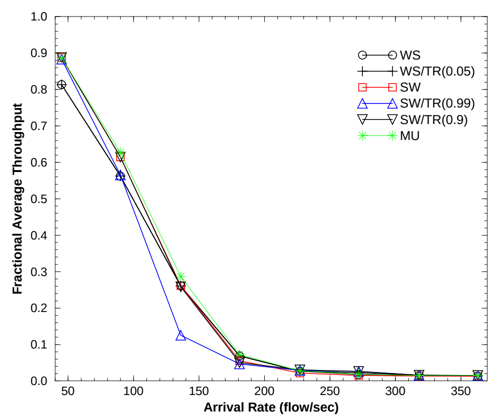

Fig. 9. Throughput comparison, US topology

Table 2. Blocking rates, Europe topology

Arrival rate	MU	WS	WS/TR(.05)	WS/TR(.25)	SW/TR(.99)	SW/TR(.9)	SW
227 flows/s	1.1%	1.2%	1.3%	1.4%	1.8 %	1.3 %	1.0 %
272 flows/s	6.3%	5.9%	6.1%	6.3%	8.4 %	6.5 %	6.3 %
318 flows/s	9.7%	9.9%	10.0%	10.1%	14.6 %	11.4 %	10.0 %
364 flows/s	13.3%	13.0%	13.3 %	13.5%	20.3 %	15.5 %	13.3 %

For WS/TR, slightly higher blocking rates are obtained with trunk reservation than without (WS), together with higher throughput. In particular, WS/TR(0.25) yields at overload, twice as much throughput as WS and slightly higher blocking rates. As expected, for lower TRL values the performance of WS/TR (resp. SW/TR) moves closer to that of WS (resp. SW). We conclude that, on the Europe topology, high values of TRL, by limiting the possibility of alternative routing, limit both the number of simultaneous active flows on alternative paths (higher throughput) and the total number of admitted flows (higher blocking). As with the previous topology, we observe that the performance of SW varies widely as a function of the TRL, while the performance of WS/TR is relatively stable. We also notice that MU adapts well to the Europe topology. MU yields high throughput under light and moderate load conditions, and low blocking under heavy load.

Results relative to the US topology (Figure 9) are fairly consistent with previous observations, though the performance difference between the algorithms is the smallest observed so far. Compared to the previous topologies, the US topology is more loosely meshed. Shortest paths are, therefore, not much more economical in terms of resources than alternative routes and presumably get rapidly saturated as load increases. The use of alternative routes at a relatively early stage explains the convergence in the performance of WS, MU and SW at moderate and high load. Trunk reservation provides WS/TR with a more effective control over network resources at high load. WS/TR(0.05), for instance, yields the lowest blocking rates (14.0% at a rate of 318 flow/sec against 16.0% for WS and 15.5% for MU) and similar throughput. The US topology highlights another effect of trunk reservation when applied to SW. It appears that SW/TR can potentially yield worse results than SW in terms of both blocking and throughput for very high TRL values (e.g. 0.99). Lower TRL values (e.g. 0.9) yield comparable throughput to the other algorithms, but higher blocking persists. This is due to the fact that trunk reservation is triggered earlier on longer routes and causes most flows to be routed on shortest paths only.

5.2 Impact of a Limited Access Rate

Previous results (without no access rate limitation) show that WS and WS/TR perform consistently well at high load, but yield poorer throughput than the other algorithms at light load. This difference in performance is likely to disappear if the throughput of elastic flows is limited by their access rate[1]. We evaluate the performance of the five algorithms considering a relatively high access rate r^+ equal to 10 BUs (representing 10% of the link capacity considered in our simulations).

Figure 5 plots the fractional average throughput as a function of the load, achieved by WS, MU, SW, WS/TR with a TRL of 0.05 and SW/TR with TRL values 0.99 and 0.9, on the full-mesh topology. Table 3 gives the corresponding

[1] or by another bottleneck link outside the considered network, or indeed by the server providing the transferred document.

Table 3. Blocking rates, $r^+ = 10BUs$, Full-mesh topology

Arrival rate	MU	WS	WS/TR(.05)	SW/TR(.99)	SW/TR(.9)	SW
530 flows/s	0%	0%	0%	0 %	0%	9.3%
606 flows/s	6.8%	6.7%	2.7%	2.3 %	2.3 %	19.5%
681 flows/s	17.7%	17.7%	10.2 %	9.5 %	9.4 %	27.2%

blocking rates. We observe that, apart from SW, all other algorithms now yield the same throughput under all load conditions. On the other hand, WS/TR(0.25) continues to give the lowest blocking rates. Throughput and blocking results also show that MU and WS now yield equivalent performance. Furthermore, we observe that a limited access rate has little incidence on the performance of WS/TR and WS at high load. Its impact is all the more important as the algorithm puts more emphasis on load balancing (For instance, SW/TR(0.99) and SW/TR(0.9) now yield the same throughput).

Results on the Europe and US topologies (not presented here) confirm that WS and WS/TR perform as well as MU at light load, while WS/TR maintains its performance advantage over WS under overload.

6 Conclusion

The main intent of this paper is to evaluate the impact of trunk reservation when used for elastic flow routing. From the algorithms *Widest-Shortest Path* and *Shortest-Widest Path*, whose performance advantages and disadvantages are well known, we derived two novel algorithms that integrate a form of trunk reservation. The latter algorithms were compared to the *Maximum Utility Path* algorithm which uses a load dependent utility function for path selection.

Simulation results show that the performance of SW/TR strongly depends on the TRL value and the topological characteristics. Although SW/TR is potentially capable of efficiently handling traffic at low and high load, the choice of the appropriate trunk reservation level remains a serious concern.

WS/TR, in contrast to SW/TR, exhibits relatively stable performance. The TRL values tested in the range of 0.05 to 0.25 yield comparable performance. Although the effect of trunk reservation on elastic flow QoS may depend on topological characteristics, WS/TR generally yields higher overall performance than WS at overload. It seems that the choice of the TRL for WS/TR is less critical than in circuit-switched networks where a TRL which is set too high causes higher blocking. In networks carrying elastic traffic, if flows are unnecessarily blocked, this still benefits ongoing flows. Simulation results show that the poor performance of WS/TR is no longer an issue when elastic flows have a limited access rate, making WS/TR rather attractive. Furthermore, algorithms like WS/TR that select minimum-hop routes in the first place, and that activate trunk reservation at overload generally yield good performance for stream-like flows. Thus, in a multiservice context, the same WS/TR algorithm could be used for routing both elastic flows and stream flows.

MU, on the other hand, generally adapts well to any network configuration, but its behavior at overload seems perfectible. Under heavy traffic conditions, MU generally yields higher blocking than WS/TR. MU does not always provide the same level of protection against performance degradation in overload as trunk reservation.

Trunk reservation with a fixed trunk reservation level is not an adaptive mechanism. In our future studies we will continue the exploration of adaptive routing schemes whose parameter(s) can be automatically adjusted to account for different network topologies and changing traffic conditions.

References

1. Sara Oueslati-Boulahia and Eliane Oubagha. An Approach to Routing Elastic Flows. In *ITC 16*. Elsevier, June 1999.
2. Z. Wang and J. Crowcroft. Quality-of-Service Routing for Supporting Multimedia Applications. *IEEE JSAC*, 14(7):1288–1294, September 1996.
3. George Apostolopoulos, Roch Guerin, Sanjay Kamat, and Satish K. Tripathi. Quality of Service Based Routing : A Performance Perspective. In *Sigcomm'98*, August 1998.
4. R. Gawlick, A. Kamath, S. Plotkin, and K.G. Ramakrishnan. Routing and Admission Control in General Topology Networks. Technical report, Stanford University, 1995.
5. L. Massoulié and J.W. Roberts. Arguments in Favor of Admission Control for TCP flows. In *ITC16*. Elsevier, June 1999.
6. L. Massoulié and J.W. Roberts. Bandwidth Sharing and Admission Control for Elastic Traffic. *ITC Seminar, Yokohama*, October 1998. to appear in Telecommunication Systems.
7. Qingming Ma. *Quality-of-Service Routing in Integrated Service Networks*. PhD thesis, Carnegie Mellon University, Pittsburgh, January 1998.
8. M. Mellia C. Casetti, G. Favalessa and M. Munafo. An Adaptive Routing Agorithm for Best-Effort Traffic in Integrated-Services Networks. In *ITC 16*. Elsevier, June 1999.
9. Qingming Ma and Peter Steenkiste. Routing Traffic with Quality-of-Service Guarantees in Integrated Services Networks. In *NOSSDAV'98*, July 1998.
10. James W. Roberts. Realizing Quality of Service Guarantees in Multiservice Networks. In T. Horsegawa et al., editor, *PMCCN'97*. Chapman and Hall, 1998.
11. Sara Oueslati-Boulahia and Eliane Oubagha. A Compartive Studing of Routing Algorithms for Elastic Flows in a Multiservice Network. Technical Report DE/DAC/OAT/79, CNET-France Telecom, September 1999.
12. J.M. Akinpelu. The overload performance of engineered networks with non-hierarchical and hierarchical routing. In *ITC-10, Montreal, Canada*, 1983.
13. Ren P. Liu and Peter J. Moylan. Dynamic Trunk Reservation for Teletraffic Links. In *GLOBCOM'95*, volume 1, pages 432–436, 1995.
14. D. Mitra and R.J. Gibbens. State-dependent routing on symmetrical loss networks with trunk reservation. *IEEE Trans. Commun.*, 41(2):400–411, Febuary 1993.