DiffServ Traffic Management with MPLS

Paulo Rogério Pereira[1], Pasquale Lepera[2], Augusto Casaca[1]

[1]INESC ID Lisboa, Rua Alves Redol, 9. 1000-029 Lisboa, Portugal.
Phone: +351-213100345. Fax: +351-213145843. Email: {prbp, Augusto.Casaca}@inesc.pt

[2]Università di Pisa. Italy. Email: leperapasquale@yahoo.it

***Abstract* — This paper describes a proportional differentiated services architecture based on MPLS constraint based routing for traffic engineering. Three classes are defined: Gold, Silver and Bronze. The Gold class gets the best service and, whenever possible, the double of the bandwidth of the Silver class. The Bronze class is a best effort class that gets the unused bandwidth from the other classes. The simulation results show that the proposed mechanisms avoid priority inversions for all load patterns. Additionally, the traffic engineering MPLS mechanisms improve the total goodput, by using alternative paths.**

I. INTRODUCTION

Quality of Service (QoS) in IP networks has raised much interest in the research and application development communities in the recent years.

Recent European projects [1] have focused on offering QoS based on Differentiated Services [2] concepts. These projects use different combinations of traffic engineering functions, admission control mechanisms for resource control, service management and business models like service level agreements or policy-based networking.

The use of Multiprotocol Label Switching (MPLS) [3] provides mechanisms for engineering network traffic patterns through the available network paths that are not subject to the limitations of the different routing protocols.

A proportional differentiation mechanism [4][5] is a refinement of a differentiated services mechanism, where an adjustable and consistent differentiation between classes is provided.

In this paper we propose to use an Olympic service model with three classes: Gold, Silver and Bronze. The Gold class should get the best QoS. Whenever possible, a proportional differentiation mechanism should give to each Gold elastic flow the double of the bandwidth as to each Silver elastic flow. The Bronze class is a best effort class that gets a very small fraction of the existing bandwidth, plus the unused bandwidth from the other classes.

This paper proposes to use a dynamic bandwidth management mechanism that configures the bandwidth of each aggregated flow, whose path and resources are reserved through MPLS Constraint Based Routing [6]. Constraint Based Routing is an important instrument for traffic engineering, since it allows choosing a path optimizing one or more metrics. Generally, a metric based on the hop count and on the available bandwidth is used.

The paper is organized as follows. Section II presents the network used for the simulations. Section III describes three scenarios and the corresponding simulation results. The first scenario does not use differentiated services or traffic engineering. The second scenario uses both these techniques. The third scenario additionally introduces the dynamic bandwidth management for proportional service differentiation. Finally, section IV draws some conclusions and points directions for further work.

II. NETWORK

A core network was built with 12 nodes, as shown in figure 1. Nodes 0-11 are the core routers, whose function is to forward traffic to its destination. Nodes 12-17 are edge routers where the traffic sources and destinations are placed. To simplify the analysis, only the three edge routers on the left generate traffic to the three edge routers on the right. All the core links were configured with a 1 Mbit/s capacity and a 10 ms delay. The capacity used in the other links is higher and sufficient to support all the traffic, because we are interested in analyzing only the traffic in the core.

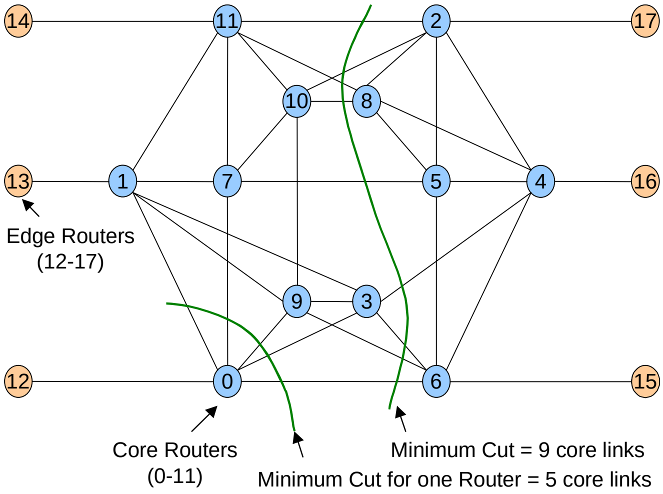

Fig. 1. Network topology.

The core should have some redundancy, so that there are alternate paths for traffic engineering. The Steiglitz method [7] was used to generate a 5-connected core network. This means that there are at least 5 link-disjoint paths between each pair of core nodes.

Work partially funded by the “EuroNGI” Network of Excelence project under IST contract 507613.

The maximum traffic the network can carry is determined by the Max-Flow Min-Cut theorem [7]. It states that the maximum flow between any two arbitrary nodes in any network cannot exceed the capacity of the minimum cut separating those two nodes. A minimum cut is defined as the set of links with minimum capacity whose removal disconnects two nodes, that is the bottleneck between those two nodes.

The maximum flow for the network core is the capacity of 9 core links, that is 9 Mbit/s, which constitutes the minimum cut from left to right, as shown in figure 1. The maximum flow an edge router can generate is 5 Mbit/s, which is the minimum cut for one core router.

The network was simulated in the NS2 network simulator [8].

III. SIMULATION RESULTS

A. First Scenario

The first scenario is a reference scenario for comparison, where no service differentiation is used, so that all traffic is treated as best effort. Additionally, no traffic engineering is used, so that all traffic is routed through the shortest path to its destination.

B. Second Scenario

The second scenario is a scenario with service differentiation with fixed bandwidths for each class. The fraction of the bandwidth of a link is as follows:

- Gold UDP: 30%
- Gold TCP: 30%
- Silver UDP: 15%
- Silver TCP: 15%
- Bronze (UDP+TCP): 5%
- Signaling traffic: 5%

These bandwidths were chosen so that there is an equal proportion of UDP and TCP bandwidth, the Gold bandwidth is the double of the Silver bandwidth and a small fraction of the bandwidth is reserved for Bronze and signaling traffic. This division makes sense if there is an equal proportion of UDP and TCP flows, and an equal proportion of Gold and Silver flows with Gold flows getting the double of the bandwidth of Silver flows.

The scheduling is performed with Class Based Queue with Weighted Round Robin (CBQ/WRR). This is a usual scheduling mechanism when several classes exist with different bandwidth reservations.

A Bandwidth Broker is used for Admission Control. It builds a Traffic Matrix of the aggregated flows for each source-destination pair, keeping track of reservations in each path. Whenever a flow exceeds the bandwidth for its class, it is downgraded to the lower available class. There is no limitation to the Bronze class that behaves like a best effort class.

Paths are selected for each aggregated flow in each class, by MPLS through Constraint Based Routing. There are 3 classes × 3 source nodes × 3 destination nodes = 27 aggregated flows. There are 27 bidirectional aggregated flows for UDP and 27 unidirectional aggregated flows for TCP. Each of the paths setup through MPLS for each aggregated flow has the fixed bandwidth fraction previously indicated. In each aggregated flow, a maximum of 10 flows were used.

The MPLS implementation used was MNSv2 [8] - MPLS for NS.

Figure 2 shows the rate per flow that each class obtains in the second scenario, when there is one Silver TCP source and a varying number of Gold TCP sources in the same aggregated flow.

These results show that, as the number of Gold TCP sources increases, the rate per flow decreases, resulting in a lower throughput for Gold flows than for Silver flows when the number of Gold flows is 3 or more. This results in a priority inversion, since the Gold flows get worse service than Silver flows. This problem is solved by the dynamic bandwidth management mechanisms proposed in the following subsection.

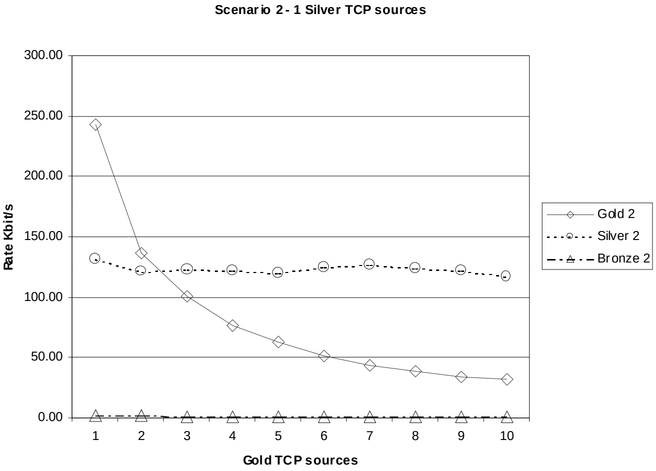

Fig. 2. Simulation results for scenario 2 - Service Differentiation with fixed bandwidths per class.

C. Third Scenario

To maintain the QoS for the Gold class consistently higher than for the Silver class, a dynamic bandwidth management mechanism is proposed. The TCP and the UDP bandwidths for the Gold and Silver classes were joined and dynamically divided as needed between the Gold and Silver classes. So, the fraction of the bandwidth of a link is now as follows:

- Gold + Silver UDP: 45%
- Gold + Silver TCP: 45%
- Bronze (UDP+TCP): 5%
- Signaling traffic: 5%

The admission control mechanism regulates the UDP bandwidth division, according to the requests of each flow, assumed non-elastic.

On the other hand, TCP flows were assumed elastic, so the bandwidth is dynamically adjusted between Gold and Silver classes, according to a proportional mechanism so that Gold class flows get the double of the bandwidth of Silver class flows.

The bandwidth is dynamically adjusted according to formulas (1)-(4). Formula (1) states that the Gold bandwidth per flow should be the double of the Silver bandwidth per flow. Formula (2) states that the total Gold bandwidth plus the total Silver bandwidth should be equal to the total available bandwidth that is 45% of the bandwidth of a core link.

$$\begin{cases} \frac{BandwidthGold}{NGoldFlows} = 2 \times \frac{BandwidthSilver}{NSilverFlows} & (1) \\ BandwidthGold + BandwidthSilver = TotalBandwidth & (2) \end{cases}$$

$$\begin{cases} BandwidthGold = \frac{2 \times TotalBandwidth \times NGoldFlows}{NSilverFlows + 2 \times NGoldFlows} & (3) \\ BandwidthSilver = TotalBandwidth - \frac{2 \times TotalBandwidth \times NGoldFlows}{NSilverFlows + 2 \times NGoldFlows} & (4) \end{cases}$$

Solving the equation system formed by equations (1) and (2), we obtain equations (3) and (4) that give the bandwidth assignments that should be used for the Gold and Silver TCP classes.

The bandwidths of the paths set up through MPLS Constraint Based Routing for each aggregated flow are configured in the third scenario according to equations (3) and (4). Only the MPLS signaling traffic is simulated in the NS simulator. The additional signaling traffic necessary for the dynamic bandwidth management is not using any bandwidth in the current simulations.

Figure 3 shows the rate per flow that each class obtains in the third scenario, when there is one Silver TCP source and a varying number of Gold TCP sources in the same aggregated flow.

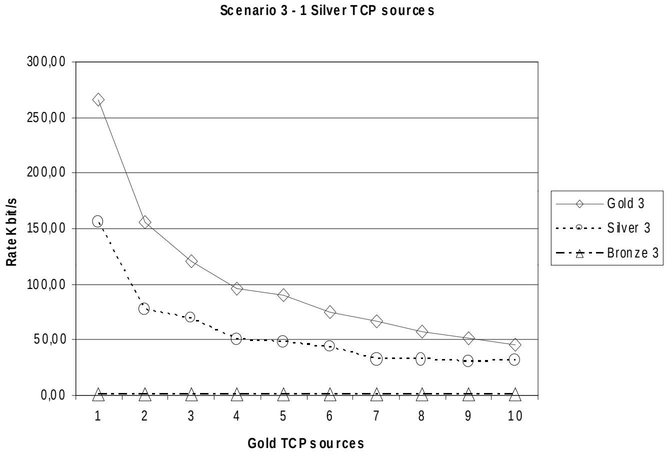

Fig. 3. Simulation results for scenario 3 – Dynamic bandwidth.

Comparing figure 3 with figure 2, the immediate conclusion is that the Gold rate per flow now is always higher than the Silver bandwidth per flow. This allows a consistent better service for the Gold class, for all traffic loads. An additional conclusion is that the Gold class gets approximately the double of the rate per flow that the Silver class gets, as expected. The result is not exactly the double because of the continuous rate adjustment of TCP sources that prevents an exact bandwidth division.

Table I shows the average TCP Goodput, which is the effective rate of user data, for the flows in each class, for all three scenarios, with the same traffic pattern in the network.

The goodput per flow shows that in the first scenario, without service differentiation, all classes get about the same goodput. In the second scenario, the bronze flows clearly get the worst service. As regards the third scenario, with dynamic bandwidth management, the Gold flows get about twice the goodput as the Silver flows, as expected.

Analyzing the total goodput in table I, it can be seen that there is about a 2.5× increase in the total TCP goodput for scenarios 2 and 3. This results from the traffic engineering mechanisms that uses alternative paths, while the first scenario only uses the shortest path, leaving some longer paths unused. This result is also shown in table II, where the total UDP goodput also shows a similar, but smaller, effect. The total goodput shown is about 8 Mbit/s, which is near the Max-Flow Min-Cut limit of 9 Mbit/s.

Table I
TCP Goodput

Rates [Kbps]	Scenario 1		Scenario 2		Scenario 3	
	Per Flow	All Flows	Per Flow	All Flows	Per Flow	All Flows
Gold	9.25	600	47.71	2793	54.51	3149
Silver	10.32	374	37.08	1537	25.58	1185
Bronze	9.50	855	3.74	337	3.86	347
Total		1829		4667		4681

Table II
Total Goodput

Goodput	Scenario 1	Scenario 2	Scenario 3
UDP	2183 Kbps	3328 Kbps	3328 Kbps
TCP	1829 Kbps	4667 Kbps	4681 Kbps
Total	4012 Kbps	7995 Kbps	8009 Kbps
Increase		1.99x	1.002x

Table III shows the UDP end to end delay. These results show that the delays are lower for the Gold and Silver classes when service differentiation is used. Additionally, the delays for all classes are lower when traffic engineering is used, in scenarios 2 and 3, since the use of alternative paths increases the total network core bandwidth that can be used.

Table III
UDP Delay

UDP Delay	Scenario 1	Scenario 2	Scenario 3
Gold	99.9 ms	30.0 ms	30.2 ms
Silver	90.8 ms	37.9 ms	37.4 ms
Bronze	92.9 ms	52.2 ms	51.8 ms

Table IV shows the UDP losses, showing that the Gold and Silver classes have no packet loss when service differentiation and admission control is used. Additionally, the packet loss rate is much higher in scenario 1, where no traffic engineering is used, since only the shortest path is used, limiting the total network core bandwidth that can be used.

Table IV
UDP Loss

UDP Loss	Scenario 1	Scenario 2	Scenario 3
Gold	35%	0	0
Silver	42%	0	0
Bronze	45%	34%	34%

The results in tables II, III and IV show that there is no significant difference in the QoS of scenarios 2 and 3, while from table I it is shown that the Gold flows get the double of the rate of Silver flows. This proves that the only effect of the dynamic bandwidth management mechanism is to keep the proportional rate differentiation between the Gold and Silver classes.

IV. CONCLUSION

The simulation results show that the use of service differentiation offers protection between classes. Excess traffic in one class does not damage the other classes.

The results also show that the use of admission control prevents UDP packet loss for Gold and Silver classes.

It is also shown that traffic engineering can improve the total goodput of the network by a factor of about 2, by using alternative paths with MPLS. It is possible to approach the Max-Flow Min-Cut theorem capacity prediction with good QoS for Gold and Silver traffic.

With the fixed bandwidth division of scenario 2, a possible priority inversion between classes can happen, e.g. with low load on Silver and high load on Gold, the Silver flows can get better QoS than the Gold flows.

The dynamic bandwidth assignment mechanism of scenario 3 assures that the priority is always respected, no matter what load pattern there is in the network. This is consistent with prices. The rate Gold flows obtain is approximately the double of the Silver flows as configured in the proportional bandwidth division. Additionally, the dynamic bandwidth division mechanism results in fewer flows having their class downgraded by the admission control mechanism. The remaining QoS parameters are similar.

As future work, other traffic patterns and other types of proportional differentiation will be tested, like proportional delay differentiation and proportional loss differentiation instead of admission control for UDP traffic.

REFERENCES

[1] S. Giordano, S. Salsano, S. Van den Berghe, G. Ventre, D. Giannakopoulos, "Advanced QoS Provisioning in IP Networks: The European Premium IP Projects", *IEEE Communications Magazine*, **41**(1):30-36, January 2003.

[2] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, "An Architecture for Differentiated Services", IETF RFC 2475, December 1998.

[3] E. Rosen, A. Viswanathan, R. Callon, "Multiprotocol Label Switching Architecture", IETF RFC 3031, January 2001.

[4] Yang Chen, Chunming Qiao, Mounir Hamdi, Danny H. K. Tsang, "Proportional Differentiation: A Scalable QoS Approach", *IEEE Communications Magazine*, **41**(6):52-58, June 2003.

[5] Nicolas Christin, Jörg Liebeherr, "A QoS Architecture for Quantitative Service Differentiation", *IEEE Communications Magazine*, **41**(6):38-45, June 2003.

[6] B. Jamoussi, L. Andersson, R. Callon, R. Dantu, L. Wu, P. Doolan, T. Worster, N. Feldman, A. Fredette, M. Girish, E. Gray, J. Heinanen, T. Kilty, A. Malis, "Constraint-Based LSP Setup using LDP", IETF RFC 3212, January 2002.

[7] Andrew S. Tanenbaum. *Computer Networks*. Prentice Hall, 1981, ISBN: 0-13-165183-8, chapter 2.

[8] UCB/LBNL/VINT Network Simulator (version 2). http://www.isi.edu/nsnam/ns/