A Case Study on Evaluating the Benefits of MPLS Traffic Engineering through Constraint-Based Routing and Network Controls[*]

S. Srivastava, B. Krithikaivasan, V. Venkatachalam, C. Beard,
D. Medhi, A. van de Liefvoort, W. Alanqar[+], A. Nagarajan[+]
School of Interdisciplinary Computing & Engineering
University of Missouri–Kansas City
Kansas City, MO 64110 USA

***Abstract*— In this paper, we attempt to understand the benefits of MPLS traffic engineering through a case study. Specifically, we do a comparative study of MPLS traffic engineering in the presence of dynamic constraint-based routing compared to destination-based routing. For this case study, we use a model derived from an actual service provider network. Different control schemes that can be deployable for MPLS traffic engineering are also considered. This case study shows that services can indeed obtain benefits from constraint-based routing and traffic engineering controls.**

I. INTRODUCTION

In the last couple of years, multi-protocol label switching (MPLS) has emerged as a viable technology for the Internet backbone. An important feature of MPLS is the ability to set up label switched paths for different services to reserve bandwidth, if and when needed. Furthermore, the possibility of doing constraint based routing in general, and for specific services, if needed, is another attractive feature.

The IETF literature (both RFCs and Internet drafts) has been deluged recently with various aspects and capabilities of MPLS [1] and the use of various features that allow deployment of controls and architectures such as virtual private networks. While most of these works describe the benefits in a qualitative manner (sometimes, from the point of view of ”good features”), very few discusses the actual quantitative benefit. For example, it had been believed for many years that the current best-effort routing (rather destination based routing) is good enough *if* enough bandwidth is available in the network. Further, there is very limited work that discuss whether different controls that can be deployed in a MPLS environment for traffic engineering are actually beneficial from the network performance standpoint.

Given this debate, we have set out to study these trade-offs. Our approach is basically simple. The team of authors for this paper, consisting of members from both academia and industry, have worked together closely to create a realistic network topology where the traffic is comprised of multiple service classes. The question we posed is: ”How would different services receive grade-of-service depending on a combination of routing possibilities and various network controls that may be placed (if MPLS is used) for traffic engineering, especially when a particular service class is required to receive better grade-of-service compared to others?” In this sense, our work is a case study. Through this case study, we attempt to gain some insight into the tradeoffs. By no means, do we address all the possible issues regarding deployment or capabilities of MPLS, such as signaling exchanges.

The rest of the paper is organized as follows. First, we discuss the different scenarios considered in our study such as destination-based routing, and different constraint-based routing schemes (with path caching). We then discuss different logical and partial sharing of link bandwidth concepts that can be deployed in a MPLS-based network for traffic engineering (note that how the actual labels may be used to deploy these concepts is outside the scope of the present paper.) In the section III, we discuss the network topology and data, and the performance metrics. The results are presented in section IV. We close with a summary of our observations for the case study in section V.

II. VARIOUS SCHEMES

In this section, we will discuss the various traffic engineering schemes used to manage networks and bring out the differences and similarities between these methods.

A. Destination Based Routing (DBR)

Destination-Based Routing uses shortest path routing based on hop count as the cost metric. This means, for a source-destination pair connected by a direct link, alternate paths will never be used to carry the traffic unless the direct link fails.

B. Constraint Based Routing (CR)

Constraint-based routing schemes attempt to find a path that satisfies one or more constraints of interest like bandwidth, delay, jitter, etc., and that is optimal with respect to some scalar metric. While a feasible path can be selected using a simple hop-count based algorithm, additional constraints can be considered to improve the resource utilization by load balancing to a certain extent. By imposing state varying constraints like bandwidth, updates are necessary for all nodes in the network in order to route incoming flows effectively. One of the major factors affecting the performance is the periodicity of these updates. If the period between

---

[*] Work supported by Sprint Corporation.

[+] Technology Planning & Integration, Sprint Corporation.

successive updates is too long, the state information is no longer valid. So, the flows can get blocked even if the best path (according to the last update) is chosen. As an offset, a crankback mechanism (used in telephone networks) can be supported by the routing mechanism to improve the performance. Since crankback involves signalling overhead, the maximum number of crankbacks should be limited to one or two.

The constraint routing schemes that we used in this study are Dynamic Random Routing (DRR), Maximum Available Capacity Routing with Periodic updates and Crankback (MACRPC), Maximum Available Capacity Routing with Periodic updates and No Crankback (MACRPNC), and Maximum Available Capacity Routing with Instantaneous Computation (MACRIC). All of these routing schemes are based on the routing computation framework discussed in [2]. DRR routes the flow by choosing a path randomly and is based on Dynamic Alternate Routing [3]. MACRPC routes the flow by hunting for a best path in terms of available bandwidth and facilitates crankback. MACRPNC is similar to MACRPC in terms of choosing a path but, doesn't allow crankback. Both MACRPC and MACRPNC use the link state information from the last updates. MACRIC uses instantaneous updates to route an incoming flow. By virtue of its nature, MACRIC is of theoretical interest. A more detailed discussion of these routing schemes can be found in [2], [4].

C. *Traffic Control*

1) Trunk Reservation (TR): Trunk reservation [5] is a simple call admission control scheme that favors direct traffic. When the available bandwidth on a link falls below a particular threshold value, the alternately routed flows will be blocked even if the flow can be accommodated. The threshold is set as a percentage of link bandwidth.

2) Service Class based Trunk Reservation (SCTR): In the previous case, all the direct flows are treated equally in accessing the TR area. In the SCTR approach, only flows of the GoS stringent service class have access to the link once the threshold is reached. This means, the non-GoS stringent class flows will be blocked similar to alternately routed flows, even if they are direct flows.

3) Service Class based Multi-Link Trunk Reservation (SMTR): SMTR extends the SCTR approach by allowing both direct and overflow traffic of the GoS stringent service class to access the TR area. This means, alternately routed flows of a GoS stringent class still can access the TR area in addition to the direct flows of a GoS stringent class.

4) Pseudo Partition (PP): Every link has a reserved area of bandwidth as a percentage of total link bandwidth for the GoS stringent service class. The arriving flows of the GoS stringent class access the reserved area first. If the reserved area cannot accommodate the new flow, the new flow is allowed to try the common area. It is a minor variant of the SMTR scheme in the order in which a flow from a GoS stringent class tries various regions for available bandwidth.

5) Partitioning (P): Due to different characteristics like delay requirements, loss requirements, different classes of service need to be treated differently by the network. The links are partitioned to construct many overlay networks based on the ratio of service class rates and their respective requirements. Each service class has its own partitioned area.

III. SIMULATION ENVIRONMENT AND NETWORK SETUP

To conduct our study, we have used the Multi-Service Dynamic Routing Simulator (MuSDyR) [7] where new control schemes such as SCTR, SMTR etc have been recently added. There is no packet level detail in this simulator which allows us to simulate thousands of simultaneous flows in an efficient manner. This allowed simulation times to be sufficiently long to produce low variance in the results over multiple simulation runs with carefully chosen seeds. The simulator supports many service models from which service classes can be constructed. The models differ in the parameters that characterize their bandwidth requirement. Some of the implemented models include Fixed Rate (FR), Uniform Fixed Rate (UFR) and Variable Bit Rate On-Off model (VBR).

The network we considered comprised of 15 nodes interconnected by 58 links. The chosen traffic can be considered as a snap-shot of traffic at one particular point in time and hence it is not representative of the performance of the network at all times. The most important aspect of the traffic model is the distribution of the load throughout the network, not necessarily the actual loading levels of the network as a whole. Among all of the source-destination pairs, 37 of them have traffic between them. Moreover, each of those selected source-destination pair has a direct link between them, although link speed might not be comparable to the offered traffic. The traffic flows are assumed to follow Poisson arrival processes with exponential holding times.

A. *Traffic Classes*

The traffic for the network model studied here comprised of four service classes namely CBR, VBR-NRT, VBR-RT and UBR with each node pair having a different parametric values for each of the service classes. Each class has its own routing table and makes its decision based on the status of the paths and the GoS requirement. The ratio between the bandwidth load product of service classes is CBR:VBR-NRT:VBR-RT:UBR:4.9:90.6:0.008:4.5. The service models used to construct these four classes are explained below:

- Constant Bit Rate (CBR) CBR is constructed using the UFR traffic model. The traffic model specifies a single bandwidth for a CBR flow, when the expectation is that flows have some variation in their bandwidth requirements. Therefore, the bandwidth of a particular CBR flow is determined from a uniformly distributed sample between 75% and 100% of the specified bandwidth. The sampling is done every time a new CBR flow arrives for a nodepair. The average flow duration is assumed as 300 secs. The mean inter arrival time is calculated based on the Erlang load and mean flow duration. Since the bandwidth is obtained by averaging over multiple connections between the given nodepair, we account for variability in the bandwidth of connections by uniformly distributing the badwidth per flow.

- Unspecified Bit Rate (UBR): UBR is implemented to have a small guaranteed bit rate, plus unspecified requirements above that. UBR is constructed from the FR traffic model. It comes as a request with a minimum CBR which is allocated to the connection and a variable rate part, for which there are no guarantees and, hence, no reservation of resources. Every active source-destination pair generates UBR traffic with a CBR part that has an inter arrival time of 10 secs and flow duration of 180 secs.
- Variable Bit Rate-non Real Time (VBR-nRT): VBR-nRT is constructed from the VBR traffic model. Sustained Cell Rates and Peak Cell Rates are provided with other parameters such as Active Burst Length assumed as 1 sec, Buffer Size taken as PCR and Cell Loss Ratio taken as 0.1%. The same bandwidth is used for all the flows for a nodepair, since all are assumed to have the same traffic parameters. The average flow duration is assumed as 600 secs. The mean inter arrival time is calculated based on the Erlang load and mean flow duration.
- Variable Bit Rate-Real Time (VBR-RT): VBR-RT is implemented the same as VBR-nRT with Cell Loss Ratio as 0.01% and flow duration as 180 secs.

B. Performance Metric

- Bandwidth Denial Ratio: It reflects the fraction of bandwidth requests that were blocked by the network of a specific service class or the overall network. Let N be the set of flows that arrive to a network and let the bandwidth requirement of flow j be given by w_j. Let N_D be the set of flows which were denied service to the network. BDR is then given by

$$BDR = \frac{\sum_{j \subset N_D} w_j}{\sum_{j \subset N} w_j} \tag{1}$$

- K% Capacity: It signifies the percentage of the baseline network capacity (100%) used in the simulation. Unless otherwise stated, it increases each link bandwidth in the same ratio.

IV. RESULTS AND DISCUSSION

We have attempted to compare routing along with traffic control mechanisms in terms of their capability to guarantee a 1% BDR for the *Constant Bit Rate* service class. The class of such mechanisms have been broadly classified into GoS passive mechanisms and GoS proactive mechanisms based on their fundamental approach. We discuss them as two scenarios here. These mechanisms facilitate a network designer to provide GoS guarantees in an effective way without having to do massive upgrades.

A. Scenario I

In this section, we experiment with the GoS passive approaches which are not biased towards any service class but intend to improve the overall network level performance. These methods have been popular in the class-less internet and the performance of a service class depends on its own traffic characterization and the overall network load. Observe that the approach is optimal co-existence rather than to reduce coupling between the classes. We want the service classes to compete with each other and help the system perform at its best.

For DBR, since many of the links have no traffic, we add capacity to the relevant links while keeping the overall addition the same as in other cases. So for an addition of $K\%$ on each link (in other cases), we do a $K_{DBR}\%$ addition on the links with direct traffic between them. Let L be the set of all the links in the network, with capacity of each link j being denoted by c_j. Let L_A be the set of links with direct traffic between them. K_{DBR} is then given by

$$K_{DBR} = \frac{K * \sum_{j \subset L} c_j}{\sum_{j \subset L_A} c_j} \tag{2}$$

For the Constraint Based Routing schemes, we add capacity to every link since the routing scheme itself is supposed to move the traffic around. We present results for the DBR and for all the four routing schemes, the BDR as seen by the CBR service class in figure 1 and by the entire network in figure 2 for increasing network capacity. Observe that the BDR of

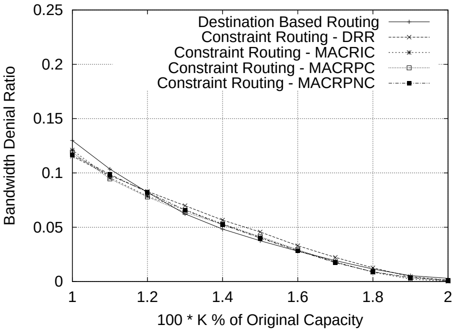

Fig. 1. BDR of CBR service class with DBR and CR schemes for increasing capacity

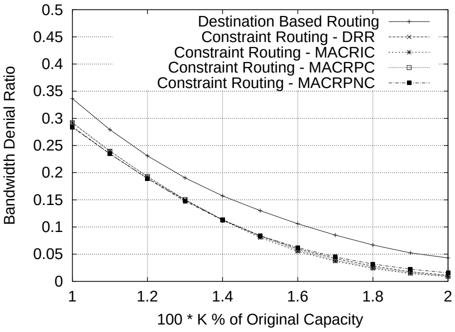

Fig. 2. BDR of Network with DBR snd CR schemes for increasing capacity

the CBR service class falls dramatically since the bandwidth of active links increases by K_{DBR}. With almost twice the present capacity in the network, the objective of providing 1% blocking can be achieved with the DBR scheme. No upgrade in the routers is needed but, the network capacity has to be nearly doubled. Moreover, it is observed that if traffic shifts with the time of day, adding selectively to specific active links will not suffice: add capacity to every link. To bring down the BDR of the entire network to 1%, the network capacity has to be increased to more than 2.6 times of its original capacity. The network requires heavy upgrade in terms of bandwidth in case an upgrade of the routers is to be avoided.

The Constraint Based Routing schemes bring down the network denial drastically from 0.35 at 100% capacity for DBR to about 0.28 at the same capacity. The benefits are reaped from the capability of the routing schemes to utilize the indirect links and help flows go across on multi-hop paths. The performance of various routing schemes does not seem to vary much and the BDR is reasonably steady for changes in routing schemes. However, the network BDR for best routing schemes does switch around 140% capacity. Hence depending upon the operating point we might choose different routing schemes. We may observe that the BDR with most routing schemes are almost equal; the one which requires minimal upgrade in the routers could be chosen. Since the routing schemes are not giving much differing performance, for here on we give the results for only two routing schemes, namely MACRPC and DRR, representatives of classes of schemes with and without periodic updates respectively.

In order to judge the impact of trunk reservation, it is incorporated on each link which is shared between multiple classes of service. Trunk reservation prevents the network from going into overloaded steady state and keeps the path lengths as small as possible. This makes the network less loaded and thereby brings down the overall bandwidth denial ratio of the network as well as that of the CBR service class. The results in terms of BDR for CBR service class are given in figure 3 and for the overall network in figure 4 with 5% and 10% trunk reservation on each link and for the above observed routing schemes. Observe that the BDR of the CBR class and the overall network remains the same as that of the Constraint Based Routing scheme implying that trunk reservation does not significantly impact the given network. We conjecture that, in a network where only a subset of all pairs have traffic between them, and the network is fully connected for this subset with additional links, lack of benefit of TR may be possible.

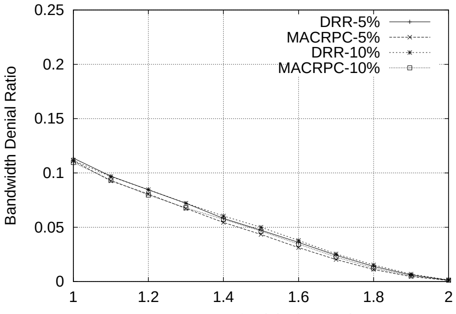

Fig. 3. BDR of CBR service class with increasing capacity for 5% and 10% trunk reservation

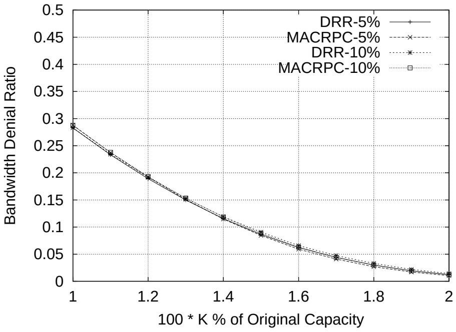

Fig. 4. BDR of network with increasing capacity for 5% and 10% trunk reservation

B. Scenario II

In the previous section, we tried to minimize the overall network BDR without providing any exclusive treatment for any particular service class and the performance of each service class depends on their traffic characteristics. The only way a provider can achieve 1% BDR for the CBR service class is by reprovisioning the entire network and bringing the overall network BDR low enough to provide the CBR service class its desired performance. In this section, we are attempting to provide special treatment to the CBR service class. There might be scenarios where the service provider intends to avoid upgrades but still aims to provide acceptable GoS to a specific service class. These approaches are very helpful in those scenarios. Of course, the network as a whole suffers more compared to GoS passive mechanisms. The fundamental idea is to weaken the coupling between the CBR and other service classes and let CBR have enough bandwidth to get its own GoS while letting other classes have relatively degraded service. We have chosen 5% and 10% (in addition, 15%, 20% in case of partitioning) bandwidth of each link and tried various techniques to see the impact on the network BDR and on the BDR of the CBR service class.

Each active source-destination pair has a direct link between them and hence we attempt to give priority to the CBR service class on the direct link which helps the network provide better performance for the CBR service class. Since the overflow from the common area flows to the trunk reserved area, the CBR service class gets direct priority over other classes. Observe that, as compared to the previous section, the coupling between the performance of the CBR service class and the other classes is weakened by reserving a section of the direct link exclusively for it and letting the flow first compete with other service classes before using the trunk reserved area. So except for the direct link, all the service classes receive the same treatment. Using this approach, the service class performance still depends upon the load offered by other classes. Hence, the CBR service class is still vulnerable to performance deterioration. We present results of BDR of the CBR service class in figure 5 and that of network in figure 6, both having a Service Class based Trunk Reservation of 5% and 10%. With 5% SCTR, the BDR of CBR service class is 5% at 100% capacity while the network denial rate is at 30%. For 10% SCTR, these numbers are 2% and 33% respectively. Hence, when bringing the BDR of the CBR service down from 5% to 2%, the network performance degrades from 30% to 33%. If the BDR of the CBR service class is to decrease to 1% using the SCTR mechanism, the returns are diminishing. To ensure 1% BDR for CBR service class, we might have to resort to partial upgrade of the link capacities and suffer a partial degrading of the overall network denial by increasing the percentage of reservation.

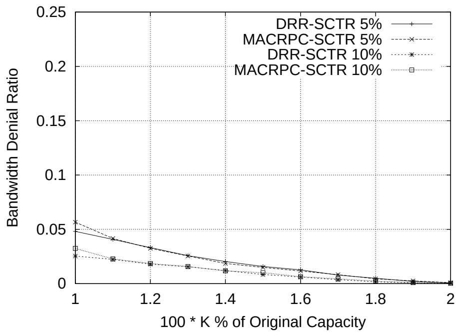

Fig. 5. BDR of CBR service class with increasing capacity for 5% and 10% service class based trunk reservation

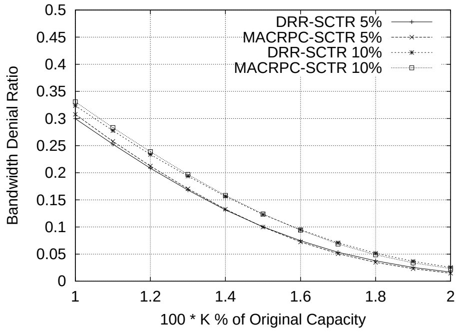

Fig. 6. BDR of Network with increasing capacity for 5% and 10% service class based trunk reservation

In the case of Service Class based Multi-Link Trunk Reservation, we let the CBR service class have a reserved area in every link which can only be used by the CBR service class either direct or overflow traffic. In this way, the CBR service class is given prioritized handling even on the indirect links. This ensures that the overflow from the direct link first tries the common area of indirect links where it competes with other service classes and if blocked, it still has a reserved area to rely on. We present results of BDR of the CBR service class in figure 7 and that of network in figure 8 for Multi-Link Service Class based Trunk Reservation of 5% and 10%.

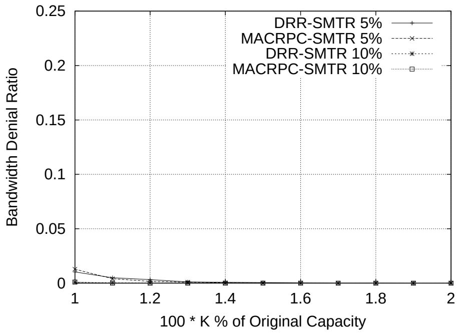

Fig. 7. BDR of CBR service class with increasing capacity for 5% and 10% multi link service class based exclusive trunk reservation

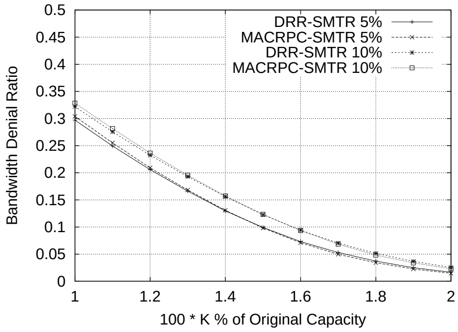

Fig. 8. BDR of Network with increasing capacity for 5% and 10% multi link service class based exclusive trunk reservation

Observe that even for 5% SMTR, the BDR of the CBR service class is nearly 1% and the overall network BDR is little less than 30% for 100% capacity. The BDR of the CBR service class can be brought into the acceptable region without any degradation in network performance. However, it may be worth observing that still the performance of the CBR service class depends on the load of other service classes. As the flows of the CBR service class always try the common area first in both direct as well as indirect paths, a fraction of flows that were accepted have taken the bandwidth from the common area. The availability of bandwidth in the common area depends on the load and BDR of other service classes. In the case of 10% reservation, the BDR of CBR service class is essentially 0% but the overall network level BDR is 33% for the same 100% capacity.

In the case of Pseudo Partitioning, the CBR service class has a reserved area in every link which can only be used by the CBR service class either direct or overflow traffic. The only difference between SMTR and pseudo partitioning is that in this case a flow of the CBR service class first tries its reserved area on the direct link and then if blocked, tries the common area of the direct link, and behaves the same way on the alternate path too. We present results of BDR of the CBR service class in figure 9 and that of network in figure 10 for pseudo partitioning with 5% and 10% reservation. As compared to SMTR, pseudo partitioning gives inferior performance for the CBR service class whereas it performs better at the overall network level. CBR service class has a minimum guaranteed throughput in both SMTR and PP mechanisms, but with PP, the CBR service class may not be able to get much more when it is ramping up due to a heavy load

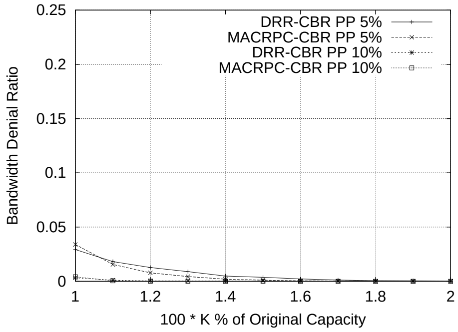

Fig. 9. BDR of CBR service class with increasing capacity for 5% and 10% pseudo partitioning

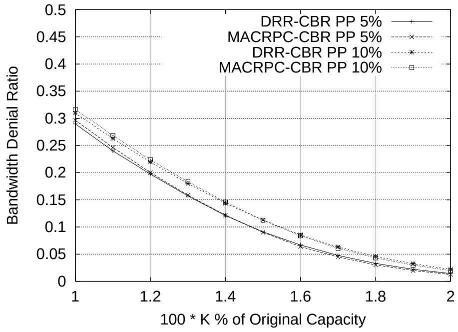

Fig. 10. BDR of Network with increasing capacity for 5% and 10% pseudo partitioning

from some other class. So, indeed, in SMTR, the reserved space is held in reserve to handle the CBR service class' peak behaviour, whereas in PP, the common area is shielded from occasional activity in the CBR service class.

In the case of partitioning, we split the network into two overlay networks; one for the CBR service class and the other one for the remaining service classes. Partitioning ensures complete immunity for GoS stringent CBR service classes from the other service classes and makes place for hard guarantees for the CBR service class. As long as the characteristics and the load of the CBR service class do not change, the performance of the CBR service class is independent of changes in the loads of other service classes. We did extensive simulation with 5%, 10%, 15%, 20% partition for CBR service class. We found that the performance is inferior compared to other controls at the same level of reservation. Hence to obtain the 1% GoS by partitioning, we need to reserve significantly more than what we need for other controls like SCTR, SMTR, etc. Also, by reserving more bandwidth, overall network performance degrades to a greater extent.

V. CONCLUSION

In this paper, we have discussed various routing and traffic engineering mechanisms which can be deployed in a MPLS based network. The MPLS framework provides the capability to have a better control of network traffic. However, the idea of best control is quite subjective and strongly depends on the network topology and the traffic characteristics. Through this case study, we have tried to answer this question for the chosen network with the considered traffic load by expressing GoS in terms of bandwidth denial ratio.

GoS passive techniques try to improve the overall network performance and in the process, give better performance to the GoS stringent service classes. In the chosen network, CR schemes provide a 5% improvement for the overall network BDR as compared to the DBR scheme. Trunk Reservation however, fails to improve upon the performance of the network provided by CR schemes.

GoS proactive approaches intend to give better performance for the GoS stringent service class at the cost of degrading the network performance. SCTR with 5% reservation decreases the BDR of CBR service class to almost half (5%), of that achievable by GoS passive schemes with 2% increase in the overall network BDR. SMTR with 5% reservation takes the gain further to 1% BDR for the CBR service class with the same 30% overall network level BDR achievable by SCTR. PP leads to slight increase in the BDR for the CBR service class while bringing down the overall network BDR as compared to SMTR. As compared to SMTR and PP, Partitioning provides degraded performance for the benefit of complete immunity. It is worthwhile to observe that providing GoS for a traffic class hinges on the degree of dependency it has with other traffic classes.

From the results we believe that for the network under study with the chosen traffic load, 5% SMTR can be considered as one of the best ways of doing traffic engineering.

We would like to draw the reader's attention towards the fact that these conclusions are only applicable and make sense with regard to the network under study. Would our results hold for other networks too? This remains to be seen as we are investigating further in this area. In fact, we don't make any such claim based on our current study. Our attempt here is primarily to report and provide some insights into the benefits and tradeoffs through a specific case study. They don't indicate in general superiority or inferiority of routing and control mechanisms with respect to one another.

REFERENCES

[1] E. Rosen, A. Viswanathan, R. Callon, "Multiprotocol Label Switching Architecture," Internet *RFC 3031*, January 2001.
[2] D. Medhi, "A Framework for QoS Routing Computation with Path Caching," Technical Report, School of Interdisciplinary Computing and Engineering, University of Missouri-Kansas City, July 2001.
[3] G. R. Ash *Dynamic Routing in Telecommunications Network*, MacGraw Hill 1998.
[4] R. J. Sivasankar, S. Ramam, S.P. Subramaniam, T. Srinivasa Rao, D. Medhi, "Some Studies on the Impact of Dynamic Traffic in a QoS-based Dynamic Routing Environment," *Proc. of IEEE ICC'2000*, pp. 959-963, June 2000.
[5] J. H. Weber, "A Simulation Study of Routing and Control in Communication Networks,"" *Bell Sys. Tech. J*, Vol. 43, pp. 2639-2676, 1964.
[6] R. Guerin and H. Ahmadi and M. Naghsineh, "Equivalent capacity and its application to bandwidth allocation in high-speed networks," *IEEE J. select. Areas Commun.*,, Vol 9, No.7, 1991.
[7] MuSDyR Manual, Technical Report, School of Interdisciplinary Computing Engineering, University of Missouri-Kansas City.