Controlling Performance in the Congested Parts of an Asymmetric Network using Controls and Routing

Balaji Krithikaivasan, Shekhar Srivastava, Cory Beard, Appie van de Liefvoort, Deep Medhi
Division of Computer Science and Electrical Engineering
School of Interdisciplinary Computing and Engineering
University of Missouri–Kansas City
Kansas City, MO 64110 USA

Abstract

In this paper, we attempt to improve the performance in congested parts (CP) of an asymmetric network using network controls and routing mechanisms. Typically, an asymmetric network is characterized by non-uniform link capacities or uneven traffic load or both. In such a network, certain service classes can experience poor performance in parts of the network, even when the average performance is very close to the expected performance level. We studied a network derived from an actual service provider network with asymmetric concentration of traffic. We have developed two metrics in order to capture the congested part behavior in a network. In order to understand the generality of our conclusions, we used three variations of our network and show that controls like Service class based Multi-link Dynamic Capacity Reservation (SMDCR) with a properly chosen reservation factor, significantly reduce variability in the performance over all parts of the network.

1. Introduction

Quality of Service routing has received a lot of attention in recent years where researchers have addressed algorithm development, performance impact and so on. In regard to performance issues, most work has so far concentrated on reporting average network behavior using symmetric networks/traffic. While average network behavior is a good metric, it does not appear to sufficiently capture network health when the network is asymmetric (capacity or traffic load or both). In other words, the average network behavior tends to give a perception that the network is doing well while some parts of the network are still congested. (We made this observation as we were doing network simulation.) Thus, our interest here is to find out the performance of the congested part of the network.

Before we delve further into this issue, we briefly discuss the role of QoS routing and network control (QoS routing and network control are discussed in detail in the next section). In separate work, it has been reported that in a QoS routing framework, activation of various network controls can play a significant role and in fact, help improve network performance that cannot be addressed by routing alone [1], [6], [9], [11], [3]. The interaction of QoS routing and network control becomes even more of an issue where one of the goals is to provide better network resource availability to a particular traffic class. (The need for better QoS arises when we address the need of higher paying customer and/or for emergency services.) Thus, to rephrase our intent, our interest here is not just limited to finding out the performance of the congested part of the network, but also to see how it can be improved in a proper QoS routing and network control environment.

In order to understand this aspect in depth, we set out to perform a case study based on data from an actual network. We have found that real networks are often asymmetric both in capacity and traffic which can be explained by asymmetric growth in traffic demands and legacy issues. Our study presented here is thus based on data obtained from a major US carrier; this model has concentration of traffic in certain areas which affects the performance significantly. In fact, our preliminary investigation (rather naive) indicated that part of network can be really badly congested.

In our study, we use a three dimensional space that addresses QoS routing with network control [11], [10]. Simply put, these three dimensions can be thought of as comprising of path caching, routing schemes and network control; they are addressed in the QoS routing framework presented in [11]. Further discussion is presented in the next section. Another important aspect is to identify the ”badly congested” or worst performing parts of the network. We have developed two metrics to identify the ”worst-part” in a network; this is discussed further in section 3.3. We present our results using the simulation tool Multi-Service Dynamic Routing simulator (MuSDyR) [7].

The rest of the paper is organized as follows. First, we review the three phase framework for constraint-based routing and discuss mechanisms for providing priority service to particular service classes. In section 3, we discuss the network topology and traffic matrices, and the performance metrics. The results are presented in section 4 for the various network models that we studied. Finally, we close with a summary of our observations.

2. Priority Mechanisms

In this work, we adopt the routing computation framework discussed in [11] in order to provide a QoS-aware environment. The framework comprises of three phases. In the first phase, a set of shortest paths are computed based on simple hop count and cached for each service class by every source to all possible destinations on the network; this is referred to as the Preliminary Path Caching (PPC) phase. In the second phase, the cached paths are ordered from most acceptable to least acceptable path using a specific routing scheme [4]; this is called the Updated Path Ordering (UPO) phase. One of the ordering criteria can be the maximal residual bandwidth along a path. The various routing schemes that we used in our study are DRR, MACRPC, MACRPNC, MACRIC along with the Destination-Based Routing (DBR) that replicates the default routing in the Internet (see [9], [10], [11] for details). In the third phase, a specific route is selected from the ordered set of paths in order to accommodate a newly arrived flow; this is the Actual Route Selection (ARS) phase. Eventually, the specific route chosen depends on the class of the flow since different controls need to be exercised depending on the service class.

We see that priority as seen by a service class is affected by the choice made in each of the three phases. As an example, the resources allocated for a specific service class in the PPC phase changes the number of routes available during the ordering phase and hence the reachable links. So, effective priority can be provided to a service class in a network by choosing a point in a three dimensional space spanned by the following dimensions, namely number of cached paths (PPC phase), type of routing scheme (UPO phase) and degree of control (ARS phase). In other words, these three dimensions completely characterize the overall prioritization mechanism being applied to a service class. The controls that are considered in the ARS phase are Dynamic Capacity Reservation (DCR), Service Class based Dynamic Capacity Reservation (SCDCR) and Service class based Multi-link Dynamic Capacity Reservation (SMDCR). DCR is a simple call admission control scheme that favors direct traffic when the available bandwidth on a link falls below a particular threshold value. SCDCR extends DCR by allowing only the direct traffic of GoS stringent flows and not all the direct flows. SMDCR further extends SCDCR by allowing both the direct and overflow traffic (multi-link traffic) of GoS stringent flows. For comparison, we also considered the case where there is no control deployed in the network which we refer in this work as No Control (NC). For more details on the controls, the reader is advised to see [8].

3. Simulation Environment and network Setup

In this work, we used the Multi-Service Dynamic Routing simulator (MuSDyR) [7]. There is no packet level detail in this simulator which allows us to simulate thousands of simultaneous flows in an efficient manner. In addition, we kept the simulation times to be sufficiently long to produce low variance in the results over multiple simulation runs with carefully chosen seeds. The flows are assumed to follow Poisson arrival processes with exponential holding times, their mean rates depending on their traffic classes. Some of the traffic models that are used to construct the traffic classes are Fixed Rate (FR), Uniform Fixed Rate (UFR) and Variable Rate on-off model (VR). UFR model is a variant of FR model where the rate is uniformly sampled in the $[\frac{3}{4}.bw, 1.bw]$ interval with bw being the effective bandwidth. For details on these models, refer [8]. Next, we describe briefly the network topology and the traffic matrices used in our study.

3.1. Network Topology

The primary network, which we call *Network I* comprises of 15 nodes connected by 58 links. This network was derived from an actual service provider network. Due to the space constraints and the nature of the topology, we do not provide a graphical picture of the topology (see [8]). The second network we study, which we call *Network II*, is derived from Network I by removing three nodes and 21 links. The links we eliminated from the network are the ones which do not carry any direct traffic and the eliminated nodes do not generate any traffic. The next section explains this further.

3.2. Network Traffic

The network traffic is comprised of four service classes. For simplicity, we call them as S1, S2, S3 and S4 respectively. S1 is constructed from UFR traffic model. S2 service class is guaranteed to have a very small bit rate (FR), plus unspecified requirements above that. S3 and S4 are constructed using VR model.

The traffic profile consists of 37 active nodepairs (has direct traffic), all of which are directly connected in both Network I and Network II. This is significant because the traffic does not need to use multi-link paths, except when the direct link cannot accommodate the traffic (and alternate routing is used). In Network I, there are 21 links idle (links that do not have direct traffic). These are the links that have been eliminated to derive Network II. The motivation to derive Network II is to increase the coupling between the traffic of different nodepairs by removing those idle links which are merely used to route the traffic alternately. This provides a more intensive look at the effectiveness of network controls.

The load ratio among the traffic classes over the network is given by *4.69 (S1): 4.58 (S2): 90.71 (S3): 0.02 (S4)* as derived from an actual network. We refer to this traffic distribution as 5% S1. From this traffic matrix, we have generated three other traffic matrices namely 10% S1, 15% S1 and 20% S1 each representing the fraction of S1 load to the overall network load. At the same time, we maintained the total network loading constant over all the traffic matrix variations. This is achieved by moving the load appropriately from traffic class S3 to S1 in the case of 10% S1, 15% S1 and 20% S1. Note that the major portion of network traffic comes from S3 flows. Since we have considered traffic class S1 as our high priority class, these traffic matrix variations allow us to see how various mechanisms perform when S1 becomes a larger fraction of the overall network load. We would like to add that the traffic is not representative of the network performance at all times. We believe that the most important aspect of the traffic is the distribution of load throughout the topology, not necessarily the actual loading levels of the network as a whole.

3.3. Performance Metric

The primary performance measures that we observe here are denoted by Ψ and ξ. These two metrics are intended to capture the Congested Parts (CP) behavior from two different perspectives. First, we would like to define *Bandwidth Denial Ratio (BDR)* (similar to classical call blocking in a multi-service environment).

For a nodegroup i, let B_{ij} be the probability of a flow from service class j getting blocked, λ_{ij} be its arrival rate, bw_{ij} be its effective bandwidth, $\mathcal{S}$ be the set of service classes and $\mathcal{N}$ be the set of nodegroups. Then,

$$BDR = \frac{\sum_{\forall i \in \mathcal{N}} \sum_{\forall j \in \mathcal{S}} \lambda_{ij}\ B_{ij}\ bw_{ij}}{\sum_{\forall i \in \mathcal{N}} \sum_{\forall j \in \mathcal{S}} \lambda_{ij}\ bw_{ij}} \tag{1}$$

In this paper, we refer to the nodepair which has the highest BDR among all the nodepairs as the *worst* nodepair either from a service class perspective or over all the service classes depending on the context.

- Ψ_j (of service class j) is defined as the weighted average of the BDR of the worst n nodepairs where n being a parameter. Formally, let $\mathcal{W}_j$ be the set of worst n nodepairs for class j, $\mathcal{N}$ be the set of all nodepairs, λ_{ij}, bw_{ij} be defined as above, BDR_{ij} be the bandwidth denial ratio of nodepair i for service class j, then

$$\Psi_j = \frac{\sum_{\forall i \in \mathcal{W}_j} \lambda_{ij}\ BDR_{ij}\ bw_{ij}}{\sum_{\forall i \in \mathcal{N}} \lambda_{ij}\ bw_{ij}} \tag{2}$$

Similarly, with $\mathcal{W}$ being the set of worst n nodepairs (over all service classes), Ψ (over all service classes) is defined as

$$\Psi = \frac{\sum_{\forall i \in \mathcal{W}} \sum_{\forall j \in \mathcal{S}} \lambda_{ij}\ BDR_{ij}\ bw_{ij}}{\sum_{\forall i \in \mathcal{N}} \sum_{\forall j \in \mathcal{S}} \lambda_{ij}\ bw_{ij}} \tag{3}$$

The utility of average BDR over all nodepairs is very much dependent on the variability of the BDR of each nodepair from the average. In this sense, this metric can be of great help for a service provider to identify the unsatisfactory customers even if the average is on par with the expected performance. In our study, we considered n to be 5.

- ξ_j (of service class j) is defined as the fraction of arrivals experiencing less than 1% BDR weighted by their bandwidth demands. Let $\mathcal{P}_j$ be the set of nodepairs that have the BDR for service class j less than 1%, λ_{ij} and bw_{ij} be defined as above. Then,

$$\xi_j = \frac{\sum_{\forall i \in \mathcal{P}_j} \lambda_{ij}\ bw_{ij}}{\sum_{\forall i \in \mathcal{N}} \lambda_{ij}\ bw_{ij}} \tag{4}$$

Similarly, with $\mathcal{P}$ being the set of nodepairs that have the overall BDR (over all service classes) less than 10%, ξ (over all service classes) be defined as,

$$\xi = \frac{\sum_{\forall i \in \mathcal{P}} \sum_{\forall j \in \mathcal{S}} \lambda_{ij}\ bw_{ij}}{\sum_{\forall i \in \mathcal{N}} \sum_{\forall j \in \mathcal{S}} \lambda_{ij}\ bw_{ij}} \tag{5}$$

We believe that a 1% BDR can be considered as a reasonable performance guarantee for a service class that warrants special treatment. In this context, ξ_j enables us to identify the fraction of nodepairs for which the guarantees are met. In other words, ξ_j identifies the unsatisfactory customers. Without loss of generality, we chose 10% BDR as the threshold at the network level as against 1% for our targeted high priority class.

3.4. Notations

The following notation is used.

K_{s1}/K_{osvc}	:	Number of cached paths for S1 service class/other service classes
RS_{s1}/RS_{osvc}	:	Routing scheme used for S1 service class/other service classes
BDR_{s1}	:	Bandwidth Denial Ratio experienced by S1 service class

3.5. Experiment Setup

All experiments were devised to provide priority to the S1 service class by varying the control parameters in PPC, UPO and ARS phases. Of particular importance is the impact of service-specific (S1 in our case) control mechanisms on the overall network performance. As we already mentioned in section 3.2, to test the robustness of our results in addition to four network variants, we have constructed four traffic matrices. In addition to Network I and Network II, we also considered two other networks by scaling down the link capacities of Networks I and II uniformly to 66% of their base capacities. For convenience, we will call them Network I$'$ and Network II$'$ respectively; they represent heavily loaded networks. In the case of Networks I and II, all the control schemes are experimented with $p/4$ and $p/2$ reservation for the S1 service class (p being the fraction of S1 service class). For Networks I$'$ and II$'$, those numbers are p and $2p$ respectively.

4. Results and Discussion

In prior work [10], the authors have compared various priority mechanisms from the average performance perspective establishing that effective prioritization of a service class can be achieved through priority in multiple phases of constraint-based routing. We set out here to explore the ability of these priority mechanisms in controlling the extent of variability away from the average behavior. We intend to fulfill our objective by looking at the variability from (a) a service class (S1) and (b) over all service classes at the nodepair level. For brevity, we present the most interesting and insightful results here.

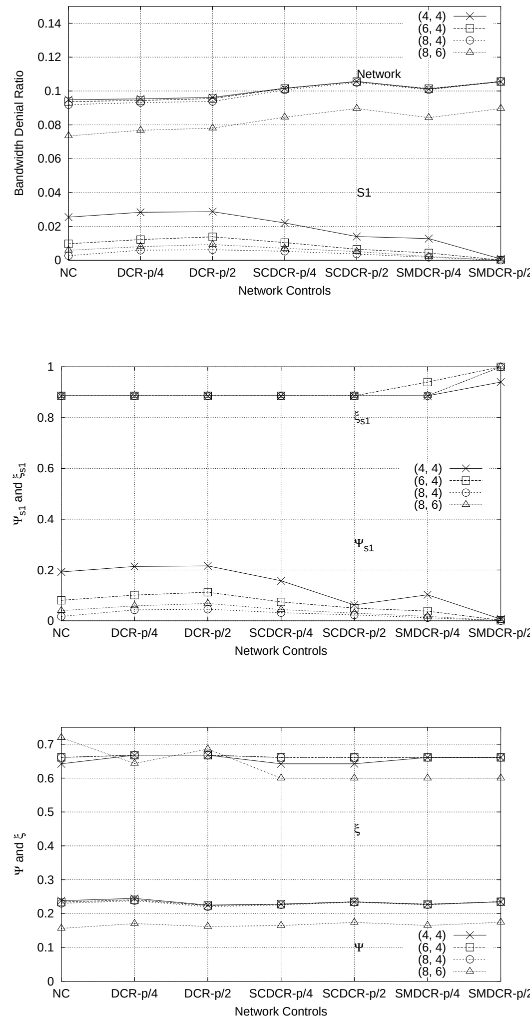

Figure 1. Network I, RS$_{s1}$ = MACRPC, RS$_{osvc}$ = DRR, 5% S1 - (a) Average Performance (b) Ψ_{s1} and ξ_{s1} (c) Ψ and ξ

4.1. Basic Results

As a first step, we would like to demonstrate the existence of high deviation from the average performance. Please note that the scale on the y-axis is not identical on the plots in Figure 1 in order to improve the clarity of them. Different curves are shown for (K$_{s1}$, K$_{osvc}$) pairs of numbers of cached paths. Figure 1(a) illustrates the average BDR of the S1 service class along with the overall network average BDR (across all service classes) whereas Figures 1(b) and 1(c) illustrate Ψ_{s1} (Ψ) and ξ_{s1} (ξ).

From Figure 1(a), it is evident that the average BDR of the S1 service class is worst for K$_{s1}$ = K$_{osvc}$ = 4 though it gets better with increasing controls. So, as we look at this case, we have the Ψ_{s1} around 20% and the average BDR around 2.5% when there are no controls (NC) in the network. We believe that such a behaviour cannot be neglected. Also, when we look at the other extreme i.e., K$_{s1}$ = 8, K$_{osvc}$ = 4, in particular with SCDCR-p/2, average BDR is close to zero. On the other hand, ξ_{s1} being around 90% (see Figure 1(b)) is indicative of the fact that still 10% of the bandwidth requests have BDR$_{s1}$ greater than 1%. At the overall performance level, except for the case K$_{s1}$ = 8, K$_{osvc}$ = 6, Ψ varies between 22% and 24% on changing controls whereas the overall average BDR stays within the range of 9.5% to 10.5% for the same (see Figures 1(a) and 1(c)). In general, service specific controls deteriorate the overall network average performance as illustrated in Figure 1(a). Such a trend is also observed in the behavior of Ψ. We believe that these results justify our attempt to look into the impact of network controls in controlling the CP behavior.

4.2. Results on Traffic Matrix Variations

In this section, we would like to discuss the impact of change in the composition of network load on the performance variability. We are more interested in the cases of RS$_{s1}$ being MACRPC or MACRIC and RS$_{osvc}$ being DRR or MACRPNC. The very basic difference between MACRPC and MACRPNC is that the latter does not facilitate crankbacks which at times can be harmful in a heavily loaded network. Moreover, choice of MACRIC can be argued on the basis of its nature of scouting for the best path for every incoming flow. Routing schemes like DRR are observed to give performance between MACRPC and MACRPNC. We discuss our observations for the two extreme traffic matrices (5% S1 and 20% S1). Due to space constraints, we show the results for the 20% S1 case.

4.2.1. RS$_{s1}$ = MACRIC, RS$_{osvc}$ = DRR. Figures 2(a) and 2(b) show the results for Network I with the 20% traffic matrix. We observed a very similar behavior of Ψ_{s1} with both 5% and 20% S1 traffic matrices. An interesting observation is that Ψ_{s1} actually increases with DCR as compared to the scenario with no controls. Also, as we look at ξ_{s1} particularly for the cases, K$_{s1}$ = 4, K$_{osvc}$ = 4 and K$_{s1}$ = 8, K$_{osvc}$ = 6, we observe a significant drop with DCR implying that more nodepairs are getting affected here. This shows that in a lightly loaded network with the load being distributed asymmetrically, uniform reservation is harmful at times. Moreover, in the case of K$_{s1}$ = 8, K$_{osvc}$ = 6, ξ_{s1} being very low indicates the high bandwidth requirement of the nodepairs being affected. Eventually, with SMDCR-$p/2$, we observe that the service level (S1) performance guarantee is met for all the nodepairs. Given this observation, we can claim that controls like SMDCR (with properly chosen fraction of reservation) not only reduces the average BDR to 1% but, also eliminates the variability significantly. The effectiveness of our claim can be understood by considering the K$_{s1}$ = 4, K$_{osvc}$ = 4 case in Figure 2(a).

If we observe Ψ from Figure 2(b), we see that the service-specific controls cause some increase in the overall CP behavior. This is because of the special treatment given to S1 traffic flows through these controls at the expense of denial of other service classes. On the other hand, in the case of the 5% traffic matrix, we observe a behavior very similar to that of RS$_{s1}$ = MACRPC, RS$_{osvc}$ = DRR. Here, the overall CP behavior does not seem to increase with the controls, which can be due to the very low load of the S1 service class with respect to other classes. In Figure 2(b), it can be seen that among all the chosen combinations of cached paths, Ψ is at its best for K$_{s1}$ = 8, K$_{osvc}$ = 6 which is very intuitive considering the increased number of cached paths for S1 as well as other service classes. However, we observe a different behavior with ξ which can be partly due

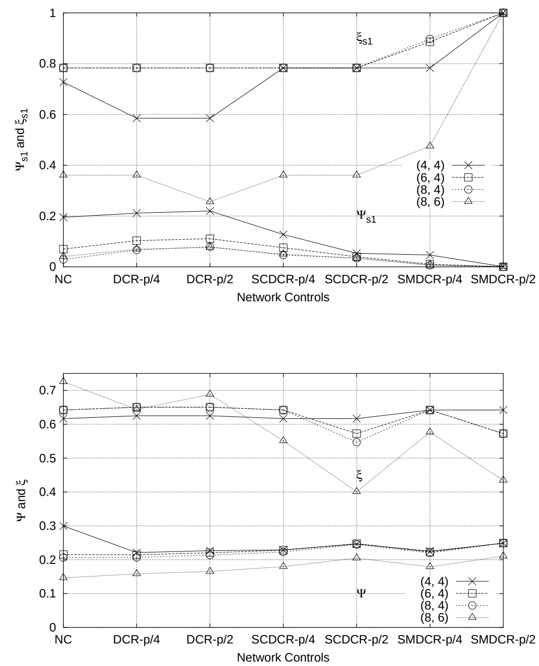

Figure 2. Network I, RS$_{s1}$ = MACRIC, RS$_{osvc}$ = DRR, 20% S1 - (a) Ψ_{s1} and ξ_{s1} (b) Ψ and ξ

to the high bandwidth nature of heavily loaded S3 flows that could have got blocked with increasing priority for the S1 service class.

4.2.2. RS$_{s1}$ = MACRPC, RS$_{osvc}$ = MACRPNC. In Figure 3, we present the results for network I with RS$_{s1}$ = MACRPC, RS$_{osvc}$ = MACRPNC. If we look at Ψ_{s1}, we made similar observations to that of RS$_{s1}$ = MACRIC, RS$_{osvc}$ = DRR. However, for the choice of K$_{s1}$ = 8, K$_{osvc}$ = 6, some improvement is seen in ξ_{s1}. We believe that MACRPC is very similar to MACRIC if the number of allowed crankbacks is same as the maximum number of cached paths.

4.3. Other Networks

In this section, we intend to observe the impact of control mechanisms on the CP behavior on three other networks which we refer to as network II, network I$'$ and network II$'$. The motivation behind such experiments is to determine the general applicability of our conclusions from the previous sections.

4.3.1. Network II. Recall here that Network II is derived from Network I with the links that do not have direct traffic

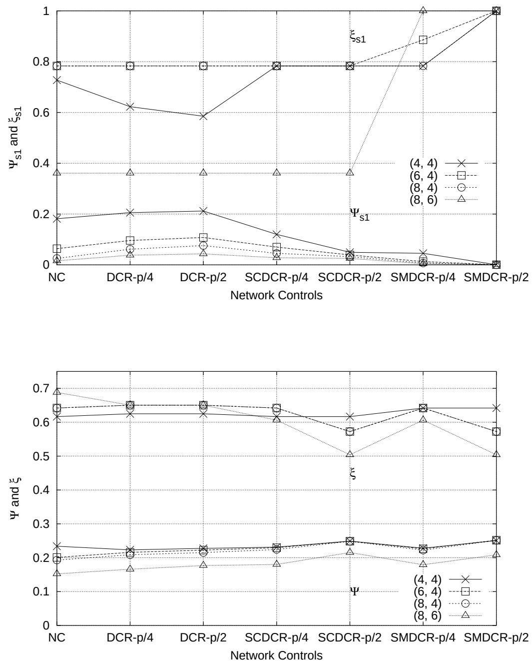

Figure 3. Network I, RS$_{s1}$ = MACRPC, RS$_{osvc}$ = MACRPNC, 20% S1 - (a) Ψ_{s1} and ξ_{s1} (b) Ψ and ξ

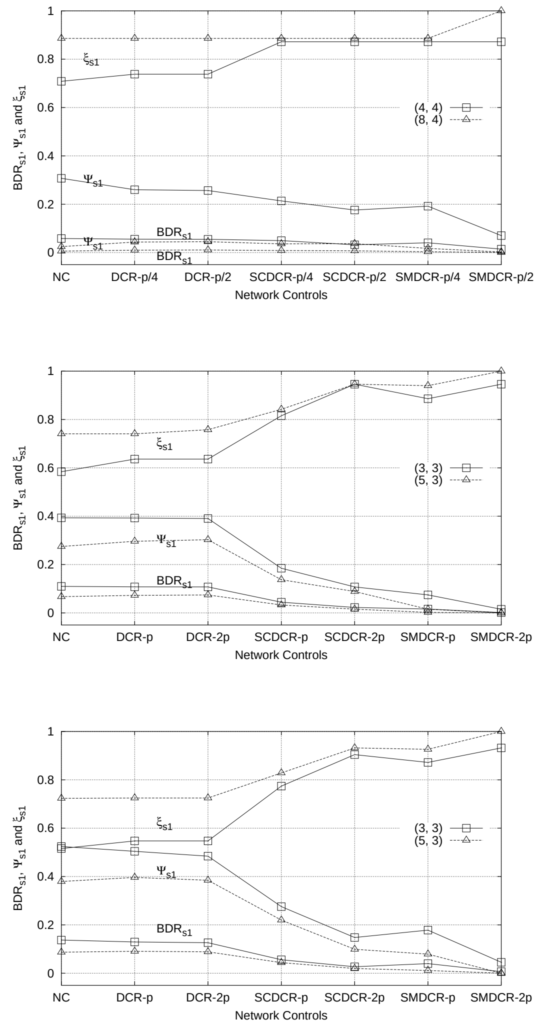

Figure 4. RS$_{s1}$ = MACRPC, RS$_{osvc}$ = DRR, 5% S1 - (a) Network II (b) Network I$'$ (c) Network II$'$

eliminated. In this context, there will be a stronger coupling between flows of different nodepairs than we have in network I. Figure 4(a) illustrates the average as well as CP behavior for the network II. We restrict ourselves here with K$_{s1}$ = 4, K$_{osvc}$ = 4 and K$_{s1}$ = 8, K$_{osvc}$ = 4 combinations in order to demonstrate the behavior at two extremes. First of all, we can look at the K$_{s1}$ = 4, K$_{osvc}$ = 4 case. We observe that Ψ_{s1} decreases significantly with DCR indicating that DCR helps here. At the same time, such implication is not very obvious for BDR$_{s1}$. Also, the further drop in Ψ_{s1} with increasing controls indicates the effectiveness of these control mechanisms very clearly as compared to the average BDR. However, as the network is heavily loaded (compared to Network I) with the number of cached paths being equal for all the service classes, Ψ_{s1} is well above 1% even with SMDCR-p/2 although, the average BDR$_{s1}$ is close to zero. In fact, ξ_{s1} also reveals such behavior. On the other hand, with K$_{s1}$ = 8, K$_{s1}$ = 4, Ψ_{s1} is closer to the average BDR which is due to the difference in the number of choices of cached paths between the S1 and other traffic classes. In summary, the impact of service specific controls on the overall BDR of the nodepairs is more significant in Network II.

4.3.2. Network I$'$ and Network II$'$. These networks are derived from Networks I and II respectively by scaling down the network capacity to 66% of the baseline capacity. Please note that p and $2p$ are the fraction of reservation used with the control mechanisms in this scenario as compared to Networks I and II. Also, we have chosen more restrictive K$_{s1}$, K$_{osvc}$ combinations considering the downsized nature of the network. If we look at Figures 4(b) and 4(c), we observe a small difference in BDR$_{s1}$ with various control mechanisms. Nevertheless, Ψ_{s1} projects a striking difference in the CP behavior bringing out the impact of elimination of idle links between Networks I$'$ and II$'$. With the SMDCR mechanism, the behavior of Ψ_{s1} in network II$'$ is very similar to that of network II (refer Figures 4(a) and 4(c)). In both these cases, S1 flows have fewer options of finding alternate routes with idle links.

In examining these three other networks, we could conclude that priority through cached paths, routing schemes and control mechanisms brings down variability to a greater extent in addition to improving the average behavior.

5. Summary

In our work, we have compared the effectiveness of various priority mechanisms that have been shown to improve the average behavior of a preferred service class significantly in [9], in controlling the deviation from the average behavior. We defined two metrics which we denoted as Ψ_j (Ψ) and ξ_j (ξ) in order to capture the CP behavior from two different perspectives (j stands for service class here). We started with a representative network and demonstrated the existence of surprising variability from the average behavior. Having done that, we observed the impact of various control mechanisms on the CP behavior for the selective combinations of cached paths and routing schemes. As an attempt to understand the applicability of our findings, we experimented with variations in traffic matrices (changing the composition of service classes). Moreover, we have experimented with three variations of our network. In this work, we have studied a sufficiently diverse set of scenarios to produce results that have general applicability.

We summarize our findings here. The CP behavior can be surprisingly bad compared to expected average performance levels. For example, the five worst blocked node-pairs could experience a blocking of 20% when the network average is merely 3% (for the 4/4 case in Figure 1) when the network is running routing algorithm MACRPC/DRR with no control scenario (NC). An approach that provides preference to priority traffic on primary paths as well as alternate paths (i.e., SMDCR) reduces variability significantly. For example, in Figure 1, the variability is *zero* with the SMDCR mechanism ($p/2$ reservation). Finally, it is important to note the applicability of providing controlled CP behavior to priority traffic for high revenue customers and emergency users. Even if emergency users could be given, say, 1% blocking overall in the network, it would also not be acceptable to have some links with 20% or more blocking. The QoS routing with network controls given here is shown to provide such CP protection. For further discussions of emergency issues, see [12].

Finally, we have shown that a point in three dimensional space of priority mechanisms can be identified for a given network and traffic matrix that can improve the service level behavior (average as well as CP behavior) without much degradation in the overall performance.

References

[1] R. S. Krupp, Stabilization of Alternate Routing Networks, *IEEE International Conference on Communications*, June 1982.

[2] F. P. Kelly, Routing and Capacity Allocation in Networks with Trunk Reservation, *Mathematics of Operations Research* 15, (1990), pp. 771-793

[3] D. Medhi and S. Guptan, Network Dimensioning and Performance of Multi-Service, Multi-Rate Loss Networks with Dynamic Routing, *IEEE/ACM Trans. on Networking* Vol. 5, no. 6, December 1997, pp. 944-957.

[4] G. Apostolopoulos, R. Guerin, S. Kamat, and S. Tripathi, Quality of Service Based Routing: A Performance Perspective, *Computer Communication Review*, Vol. 28, no. 4, October 1998.

[5] D. Medhi and I. Sukiman, Multi-Service Dynamic QoS Routing Schemes with Call Admission Control: A Comparative Study, *Journal of Network and Systems Management*, Vol. 8, no. 2, pp. 157-190, June 2000.

[6] R. J. Sivasankar, S. Ramam, S.P. Subramaniam, T. S. Rao, D. Medhi, Some Studies on the Impact of Dynamic Traffic in a QoS-based Dynamic Routing Environment, *2000 IEEE International Conference on Communications*, pp. 959-963.

[7] MuSDyR Manual, School of Interdisciplinary Computing Engineering, University of Missouri-Kansas City (available at `http://conrel.sice.umkc.edu/simulator/Musdyr1.1.pdf`).

[8] B. Krithikaivasan, S. Srivastava, C. Beard, D. Medhi, A. van de Liefvoort, Controlling Performance in the Congested Parts of an Asymmetric Network using Controls and Routing, *extended version*, (available at `http://conrel.sice.umkc.edu/bkrith/papers/iscc_ext.pdf`.

[9] S. Srivastava, B. Krithikaivasan, V. Venkatachalam, C. Beard, D. Medhi, A. van de Liefvoort, W. Alanqar, and A. Nagarajan, A Case Study on Evaluating the Benefits of MPLS Traffic Engineering through Constraint-Based Routing and Network Controls, *2002 IEEE International Conference on Communications*, April 2002.

[10] S. Srivastava, B. Krithikaivasan, C. Beard, D. Medhi, A. van de Liefvoort, W. Alanqar, A. Nagarajan, Benefits of Traffic Engineering using QoS Routing Schemes and Network Controls, *submitted to Computer Communications*.

[11] D. Medhi, Quality of Service (QoS) Routing Computation with Path Caching: A Framework and Network Performance, *IEEE Communication Magazine*, December 2002.

[12] C. Beard and V. Frost, Prioritized Resource Allocation for Stressed Networks, *IEEE/ACM Transactions on Networking*, Vol. 6, no. 5, October 2001, pp. 618-633.