Stateless Proportional Bandwidth Allocation

Prasanna K. Jagannathan[*a], Arjan Durresi[*a], Raj Jain[**b]
[a]Computer and Information Science Department, The Ohio State University

ABSTRACT

This paper proposes a new method of bandwidth allocation during congestion, called the proportional allocation of bandwidth. Traditionally, max-min fairness has been proposed to allocate bandwidth under congestion. Our allocation scheme considers the situation where flows might have different subscribed information rates, based on their origin. In proportional allocation of bandwidth, during congestion all flows get a share of available bandwidth, which is in proportion to their subscribed information rate. We suggest a method for implementing this with minimum signaling and without storing per-flow state. Our method is based on the principles of differentiated services – diffserv. We show by simulation that it is possible to obtain proportional bandwidth allocation without per-flow state storing and minimum signaling.

Keywords: Proportional allocation of bandwidth, differentiated services, layering, link-sharing

1. INTRODUCTION

The current Internet provides best effort service and relies on TCP mechanisms for end-to-end congestion control. Even though this solution has been very successful in preventing network collapse, it doesn't always guarantee fairness among users during congestion. To provide fairness in sharing the bandwidth during congestion, the max-min fairness principle has been proposed[1]. There exist many mechanisms for achieving max-min fairness using per-flow state information. However implementing these mechanisms at the routers in the core of the network where there are millions of flows might be difficult. Mechanisms like CSFQ[1], which do not store per-flow state information achieve approximate max-min fairness among competing flows. However in the current Internet, users might have different service requirements. The most important service parameter the subscribed information rate (SIR) varies widely from one user to another. However in the best effort service of the Internet, SIR is not considered while allocating bandwidth during congestion. A user with a higher SIR generally pays more than a user with a lower SIR. Therefore during congestion, a user with a higher SIR will expect to be allocated more bandwidth than the bandwidth allocated to a user with a lower SIR. SIR could be assigned based on the access rate. Other more sophisticated schemes for SIR are possible, for example the traffic priority could be taken into consideration. We define a new method of bandwidth allocation called Proportional Allocation of Bandwidth – PAB, in which bandwidth is allocated in proportion to the SIR of the competing flows.

In our technique to implement PAB, no information about the state of the flows is stored in the interior of the network. We avoid the storing of per-flow state information by encoding the ratio of a flow's data rate to its SIR in the form of a label on its packets. At the interior of the network, the routers use these labels for differentiating between packets during congestion. All the labeling is done at either the source or the first network element – the ingress router after the source, which has information about the source's SIR. The core router drops packets based on their labels and the current level of threshold in the router. Thus no state information is stored in the center of the network.

We assume a network model, which is similar to Differentiated Services[16], CSFQ[1] and RFQ[2, 3] as shown in Fig. 1. The source or the edge router marks the labels in the packets. The data rate of each flow is estimated. Using the flow's data rate and the SIR of the flow, the packets of the flow are marked with different labels. The labels indicate priority and a packet with higher priority has less probability of being dropped than a packet with a lower priority.

[*] durresi@cis.ohio-state.edu; phone: 1 614 688 5610; fax: 1 614 292 2911; http://www. cis.ohio-state.edu; Computer and Information Science Department, The Ohio State University; 2015 Neil Ave., Columbus, OH, 43210 ,USA
[**] raj@nayna.com; phone: 1 408 956 8000X309, http://www.nayna.com, Nayna. Inc.; 157 Topaz St. Milpitas, CA 95035, USA

The core router uses multilevel drop thresholds to drop packets, based on their label value. Depending on the severity of congestion, the core router varies the drop probability.

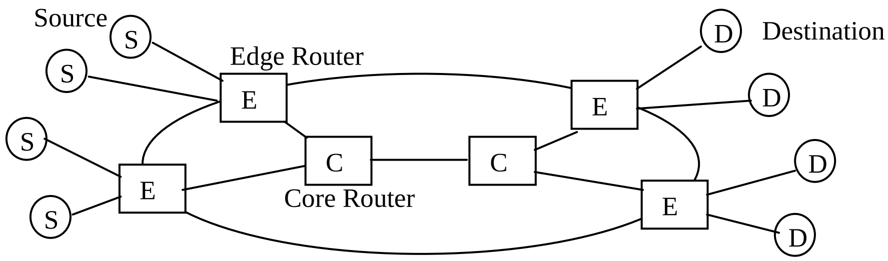

Fig. 1. Network Configuration

Simulation results show that our technique for allocating bandwidth according to PAB, has very good performance and achieves proportional bandwidth allocation without storing any state information.

The organization of the paper is as follows: Section 2 deals with the operation of PAB-Proportional Allocation of Bandwidth. Section 3 deals with technique for implementation of PAB. In Section 4 we discuss about the simulation results of our implementation. We conclude in Section 6.

2. PROPORTIONAL ALLOCATION OF BANDWIDTH - PAB

The principle behind PAB is that the allocation of bandwidth should be in proportion to SIR of the flows sharing the link. The SIR of the source of a flow is one of the most important service parameters for a flow. It is therefore important to consider both a flow's data rate and its SIR, to allocate bandwidth.

According to the definition of max-min fairness as defined in[1], each flow is allocated bandwidth as given by the Formula (1) and (2)

$$Alloc\ (i) = Min\{send(i),\ rr\} \tag{1}$$

$$\sum Alloc\ (i) \leq Available\ Bandwidth \tag{2}$$

Here *send(i)* is the data rate of the i^{th} flow and *rr* is the maximum rate that satisfies the above inequality. Any flow sending more than *rr* will have its throughput reduced to *rr*. However under this scheme, the SIR of the flow is not considered in bandwidth allocation.

We suggest that bandwidth be allocated such that all flows have identical flow rate to SIR ratio. However this requirement must be satisfied with full network utilization. Therefore in PAB the allocation of bandwidth is given by Formula (3) and (4).

$$Alloc(i) = Min\{\ send(i),\ frac * SIR(i)\ \} \tag{3}$$

$$\sum Alloc(i) \leq Available\ Bandwidth \tag{4}$$

Here, SIR(i) is the SIR of the i^{th} flow and *frac* is the maximum fractional multiplier (between 0 and 1) that satisfies the above inequality. The *frac* determines the maximum data rate of a flow as a fraction of its SIR. If the data rate of a flow is below its allowed throughput *frac*SIR* then it does not suffer any packet loss. Further if a flow has a data rate less than its allowed fraction of SIR, then the remaining excess bandwidth is also shared among other flows in proportion to their SIR. No flow is allowed to send more than its SIR during congestion. The throughput of any flow sending more than the allowed fraction of SIR is reduced to its maximum allowed data rate. Thus PAB differentiates between flows and allocates bandwidth in proportion to the SIR of the flows.

We describe an example illustrating the proportional allocation of bandwidth: Consider 3 sources, S_1, S_2 and S_3 sending one flow each to destinations D_1, D_2 and D_3 respectively. Let the flows traverse two congested links l_1 and l_2. The network configuration is as shown in Fig. 2.

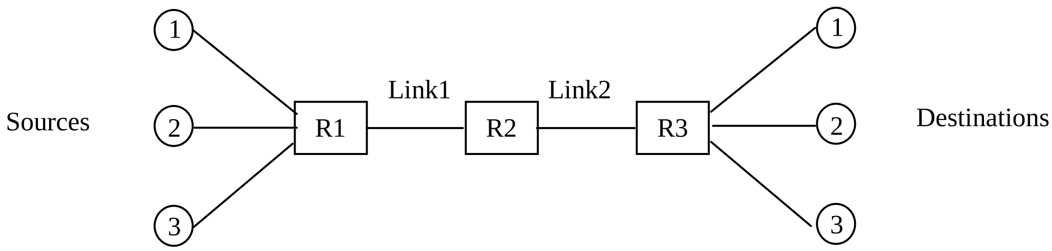

Fig. 2. Network Configuration for PAB Illustration

Let the SIR of the sources be 100 Kbps, 150 Kbps and 300 Kbps.. Let the bandwidth of the links l_1 and l_2 be 200 kbps and 150 kbps respectively. The flow rate (FR) before the links for the sources are 90 kbps, 120 kbps and 200 kbps. The throughput of the sources after l_1 and l_2 are given in TABLE I and shown in the Fig. 3.

The ratio of the SIR of the individual sources to the sum of SIRs of all sources determines the throughput of the sources during congestion. The throughput of a source is the product of the ratio of its SIR to the total SIR and the bandwidth of the bottleneck link. The sum of the throughput of the sources after link l_1 and before l_2 is 200 kbps and that after link l_2 is 150 kbps, utilizing the full capacity of the network.

TABLE I
SIR, FR and Throughput for PAB Illustration

Sources	**SIR (Kbps)**	**Ratio of SIR to ΣSIR (%)**	**FR before link l1 (Kbps)**	**FR after link l1 and before link l2 (Kbps)**	**FR after link l2 (Kbps)**
S1	100	18.18	90	36.36	27.27
S2	150	27.27	120	54.54	40.90
S3	300	54.54	200	109.09	81.81

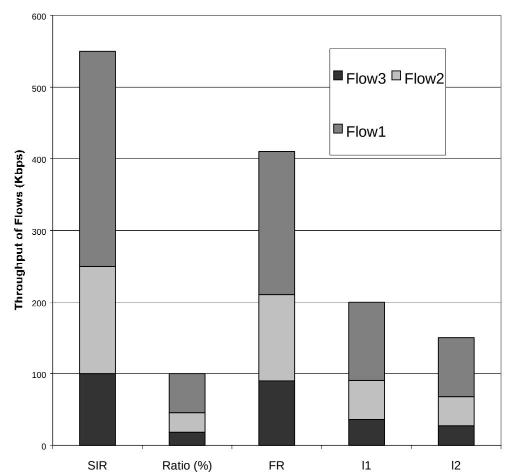

Fig. 3. PAB Ratio and Throughput

3. IMPLEMENTAION OF PAB

Our technique for implementing PAB is based on the principles of Differentiated Services[16]. To avoid per-flow state information in the core routers, our technique uses labels to indicate the ratio of flow rate to the SIR of the flow. Packets are marked with labels at the edge of the network. The ratio of flow rate to the SIR of each flow is encoded as the label on the flow's packets. In the center of the network, bandwidth allocation is done using these labels to differentiate between packets. The core router performs multilevel threshold based dropping.

Our technique to implement PAB involves two main components: the labeling of the packets at the edge of the network and dropping of packets at the core router.

3.1 Packet labeling methodology

Packets are labeled at the source or the ingress router. The ingress router has knowledge of SIRs of all the sources connected to it. The labeling mechanism marks the flow's packet with different labels depending on the ratio of the flow rate (FR) to the SIR. The total number of available labels is fixed for all flows, but the number of labels used at any time for a source depends on its FR. As the ratio of FR to SIR increases for a flow, more and more packets will be marked with labels with low priority. We describe how this mechanism marks packets depending on flow rate and the SIR of the flow rate. This mechanism marks the packets of two flows with the same FR but different SIR, differently. The flow with lower SIR has more packets with lower priority than the flow with the higher SIR as shown in Fig. 4.

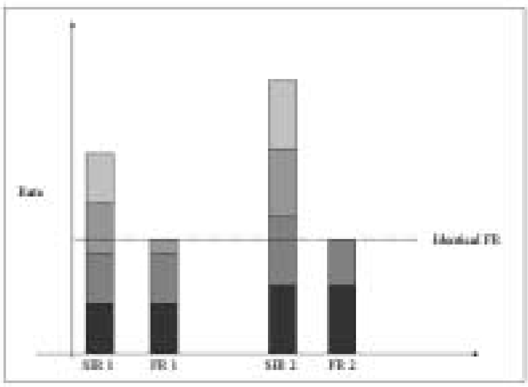

Fig. 4. Packet Labeling in flows with same FR but different SIR

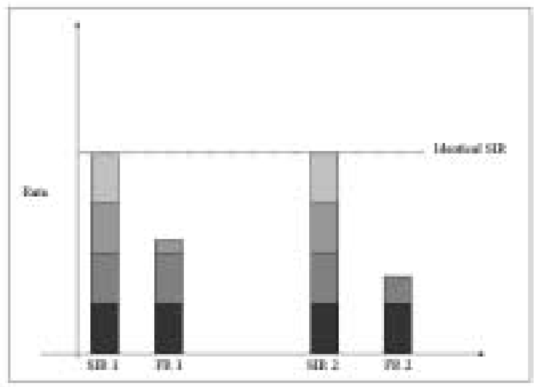

Fig. 5. Packet Labeling in flows with same SIR but different FR

This mechanism also marks the packets of two flows with the same SIR and different FR differently. The flow with the higher FR is marked with more lower priority labels than the flow with the lower FR as shown in Fig. 5. A label is not associated with a particular rate numerically. Each label is associated with only a fraction between 0 and 1. The sum of the fractions corresponding to all labels is set equal to one. In Fig. 4 and Fig. 5, we assume that there are only four labels and all labels are associated with the same fraction value ¼.

To label the packets, multiple token buckets are used. The source should be able to send data at or below its SIR, but its average data rate should not exceed its SIR. So the sum of the token rates of all the token buckets must be equal to the SIR of the source. The SIR however is distributed among all the token rates. Therefore the token rate of an individual token bucket is a fraction of the source's SIR. This fraction is equal to the fraction associated with the label corresponding to that token bucket. The sum of all the fractions associated with the labels is 1. So the sum of the token rates is the SIR of the source. The significance of the value of the fractions will be discussed later. The flow rate determines the actual label values that the packets get. As the ratio of flow rate to SIR increases, more and more packets will be labeled with lower priority. The token bucket size allows for bursts in the flow rate. However the long term rate of the flow can never exceed its SIR.

The lower the label value, higher is the priority of that packet and it has lower probability of being dropped. A packet can remove tokens from only one bucket. If a packet has insufficient tokens to remove from any of the token buckets then the source is sending packets at a rate greater than SIR and the available burst size. So the packet is dropped. The Fig. 6 and Fig. 7 show a pictorial representation of the token bucket system used for labeling packets. In Fig. 6 packet(*I*) waits to remove tokens from a token bucket. Token buckets *1* to *K-1* do not have sufficient tokens required by packet *(I)*. So the packet *(I)* removes tokens from bucket *K*. The tokens in the token bucket *K* have decreased due to the consumption of tokens by packet *(I)* as shown in Fig. 7. Since packet *(I)* removes tokens from the token bucket *K*, packet *(I)* is marked with label value *K*. The number of tokens in the token bucket K has decreased due to token consumption by packet *(I)*.

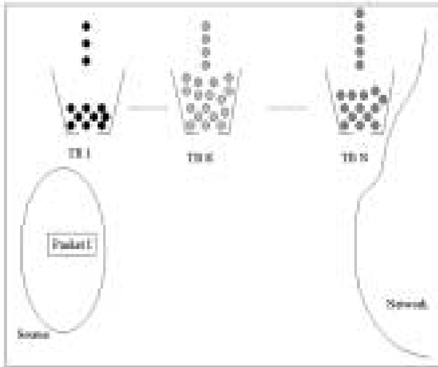

Fig. 6. Before Packet Labeling

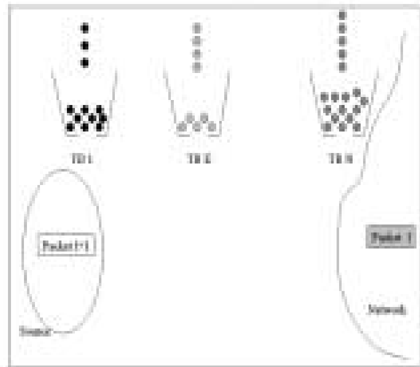

Fig. 7. After Packet Labeling

3.2 Packet Dropping Mechanism at the Core Routers

At the core router, packets are dropped based on their labels. The average queue size is monitored using exponential weighted moving average technique. As the average queue length changes packets are dropped with correspondingly changing probabilities. The drop probability for a packet with a lower priority label is higher than a packet with a higher priority label for the same average queue length. The active queuing mechanism that is used in our technique is similar to SAMT - Single Accounting Multiple Thresholds scheme of [16] and RIO scheme of [4]. It has n levels of RED with a single common average queue length.

In n-RED, multi level drop thresholds are established at the core routers and the drop probability of a packet is calculated based on its label value and the average queue length. A single average queue length is maintained for all the packets. There are n different label values and n sets of (min_{th}, max_{th}, P_{max}) exist in the core router, where min_{th} indicates the minimum threshold and max_{th} indicates the maximum threshold and P_{max} indicates the maximum drop probability. When a packet of label value k arrives in the router, the min_{th}-k and max_{th}-k and P_{max}-k are used to determine whether that packet has to be dropped. The priority among different label is achieved by choosing correct values for thresholds and drop probabilities. The values are set as shown in Fig. 8. A low priority label has its thresholds such that it has high drop probability even when the average queue length is low. However a high priority label has its thresholds such that it has a low drop probability even when the average queue length is high. The highest priority label is Label 1 and the least priority label is Label N. Lower priority labels have high maximum drop probability.

3.3 Determination of the label fractions

We have studied three different sets of fraction values that can be assigned to the labels. These three sets of fraction values have specific properties. They are:

a. Fractions with equal value - Equal fractions
b. Fractions forming arithmetic progression - AP fractions
c. Fractions forming geometric progression - GP fractions

3.3.1 Equal fractions

All the fractions are of equal value. So, if there are N labels, then each label has the value 1/N. So the sum of fractions is 1.

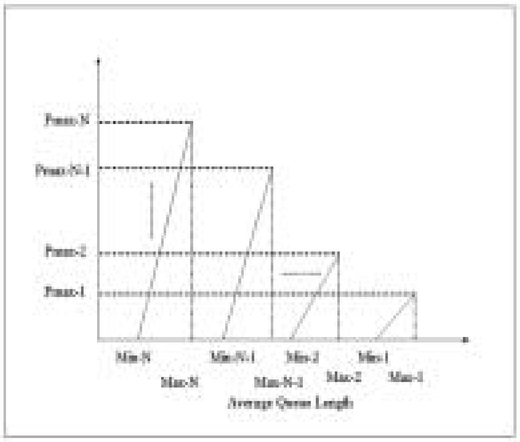

Fig. 8. Core Router Mechanism for Dropping Packets.

3.3.2 AP fractions

The fractions form an arithmetic progression. Unlike equal fractions, the values for AP fractions are not identical. To achieve better granularity while providing proportional bandwidth allocation, the smaller values in the arithmetic progression are associated with the higher priorities among the labels. So the fractions will have the values given by:

$$a,\ a+d,\ a+2d,\ ..\ ,\ a+(N-1)d \tag{5}$$

Since the sum of the fraction values in the arithmetic progression must be unity, the sum of the arithmetic progression is set equal to one.

$$Sum = Na + \frac{N(N-1)d}{2} = 1 \tag{6}$$

For simplification, we assume '*a*' to be equal to '*d*'. So this gives the values of the fractions as

$$d, 2d, 3d, .., Nd$$

We obtain the following value for '*d*' when the sum of the fractions is unity ,

$$d = \frac{2}{N*(N+1)} \tag{7}$$

When there are 8 labels, N becomes 8. After solving for N=8 in the above equation the value of d is 1/36. Therefore the values of fractions are 1/36, 2/36, 3/36, 4/36, 5/36, 6/36, 7/36, and 8/36.

3.3.3 GP fractions

The fractions can also form a geometric progression. Similar to arithmetic progression, the values of the fractions are assigned such that higher priority labels are associated with smaller values in the geometric progression. So the fractions are given by:

$$a,\ ar,\ ar^2,\ ..\ ,\ ar^{(N-1)} \tag{8}$$

Since the sum of the fractions should be 1, the value of sum is set to one.

$$Sum = \frac{a(1-r^N)}{1-r} = 1 \tag{9}$$

Again as in arithmetic progression, to simplify calculations, we assume 'a' to be equal to 'r'. So the values of the fractions are given by:

$$r,\ r^2,\ r^3,\ ..\ r^N \tag{10}$$

The value for r is given by the equation below when the sum of the fractions is unity.

$$r\ (2 - r^N) = 1 \tag{11}$$

For N = 8 , r gets a value approximately equal to ½. So the fractions have the following values : 1/2, 1/4, 1/8, 1/16, 1/32, 1/64, 1/128 and 1/128. The last two fractions are made equal so that the sum of the fractions is 1.

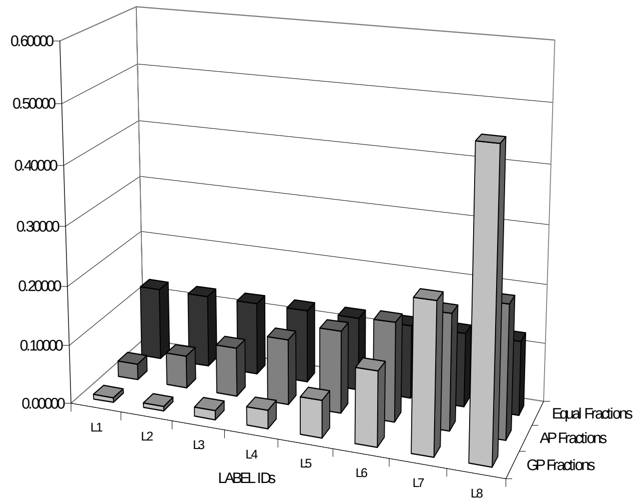

Fig. 9. Value of Equal, AP and GP Fractions

The Fig. 9 shows the values of fractions assigned to labels for the three types of fractions – EF, AP Fractions and GP Fractions. For the same SIR and the same flow rate the number of high priority packets is more in equal fractions than in AP fractions than in GP fractions. However for the same SIR and the same flow rate the number of low priority packets is less in equal fractions than in AP fractions, which is less than that in GP fractions.

4. SIMULATION RESULTS

In our simulations we compared the performance of our implementation technique with equal fractions, AP fractions, GP fractions and simple RED. We studied the performance of our technique in single congested link and multiple congested links.

4.1 Single Congested Link

We used the ns-2 simulator [18, 19] for performing simulations. The packets are marked with labels at the edge of the network at the ingress routers. The core routers use the labels for providing service and dropping packets, when there is congestion.

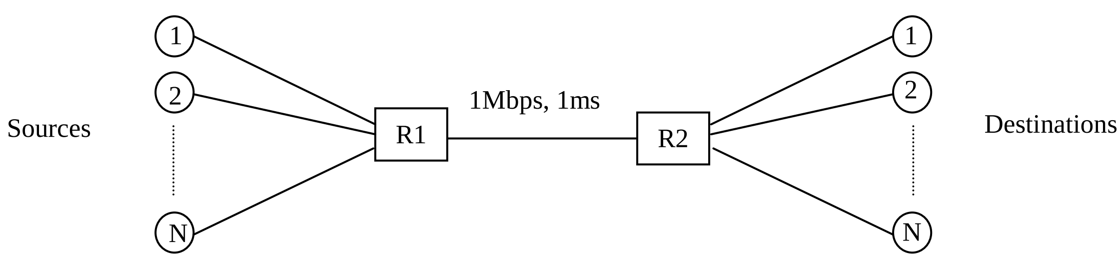

Fig. 10. Single Congested Link - Network Configuration

In the case of single congested link, the network configuration is as shown in Fig. 10. There are N flows sharing a single congested bottleneck link. The bottleneck link is at the core of the network. The SIR of the i^{th} source was set at $\left(\frac{i}{N} * 500\right)$ Kbps. The number of sources sharing the link varied from 3 to 32. The capacity of the bottleneck link is 1

Mbps and the link delay is 1 ms. The capacity of the link buffer was 100 packets. The packet size of the TCP sources was set at 1000 bytes. The packet size of the CBR sources was set at 210 bytes. The parameters for 8-level drop thresholds at the core router are shown in the TABLE II.

TABLE II. Parameter for Core Router Dropping Mechanism

Label	Minimum Threshold	Maximum Threshold	Max Drop Probability
1-highest priority	80	90	1/50
2	70	80	1/45
3	60	70	1/40
4	50	60	1/35
5	40	50	1/30
6	30	40	1/25
7	20	30	1/20
8-lowest priority	10	20	1/15

For the token buckets at the ingress routers, the bucket size was fixed at 80000 bytes. The token rate for each label was determined depending on the method of fractions chosen for marking the labels. The token rate for each of the eight labels is the product of the fractions associated with that label and the SIR of the source.

By definition of PAB, each flow should get a share of bandwidth, which is in proportion to its SIR.
The measure that we used to calculate the effectiveness of the proportional allocation of bandwidth is obtained as shown. The throughput ratio of the i^{th} flow [TR(i)] is defined as the ratio of throughput of i^{th} flow to the sum of the throughputs of all flows going through the same link.

$$TR(i) = \frac{Throughput - of - flow(i)}{\sum Throughput - of - all - flows} \tag{11}$$

The SIR ratio for the i^{th} flow [SR(i)] is defined as the ratio of SIR of i^{th} flow to the sum of the SIRs of all flows going through the same link

$$SR(i) = \frac{SIR - of - Source(i)}{\sum SIR - of - Source(i)} \tag{12}$$

The Allocation ratio for the i^{th} flow [AR(i)] is defined as the ratio of TR(i) to SR(i).

$$AR(i) = \frac{TR(i)}{SR(i)} \tag{13}$$

The performance measure is given by the following formula.

$$ProportionalityIndex = \frac{\left[\sum AR(i) * \sum AR(i)\right]}{N * \sum \left[AR(i) * AR(i)\right]} \tag{14}$$

In the first set of experiments the N sources were UDP. All the UDP sources were CBR sources. The data rates of the sources were set randomly from 10% to 200% of their SIR. The experiments were performed using Equal Fractions, AP Fractions, GP Fractions and RED. The Fig. 11 shows the performance of all three methods of fractions. For equal fractions 30 simulations were performed, with the number of sources flowing through the single congested link increasing from 3 to 32. The network configuration for the experiments of AP fractions, GP fractions and RED was similar to the network configuration of equal fractions. Similar to equal fractions, 30 experiments were done for AP fractions, GP fractions and RED.

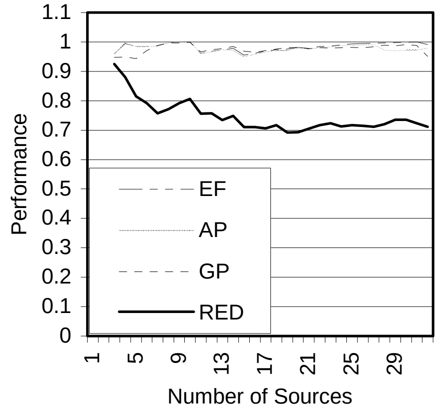

Fig. 11. Performance in Single Congested Link with UDP sources

From Fig. 11 it can be observed that as the number of sources increases, the performances of the three types of fractions are almost similar. However for RED, the performance drastically drops as the number of sources increases. This is due to the fact that RED has no knowledge of SIR and thus cannot differentiate between the flows based on their SIR and thus the bandwidth allocation by RED does not follow the principles of Proportional Allocation of Bandwidth. Actually RED was not designed to enable sharing of resources proportional to SIR. The reason we are comparing PAB with RED is to show how much would benefit users with a PAB service. In the second set of experiments the N sources were TCP. The TCP sources were Telnet applications with their peak rates set at 400 Kbps. As before the number of sources increased from 3 to 32 and the three types of fractions and RED were used to perform simulations. The Fig. 12 shows the performance for TCP sources.

In the case of TCP sources, as the number of sources increases the performance of the three fractions are very good. The TCP flows are congestion sensitive and when there is congestion, the TCP flows tend to share the bandwidth equally among the flows. Our technique achieves proportional bandwidth sharing by using labels and thus achieves good performance. In the case of RED, the performance has become much worse than that with UDP sources. During congestion the TCP sources reduce the sending data rate so that the rate of all sources are equal and RED cannot distinguish between sources and thus has poor performance.

In the third set of experiments the N sources were mixed. The sources 0,2,4,.. were TCP Telnet sources with the same parameters as before and the sources 1,3,5,.were CBR UDP sources with the same data rate as before. As before the number of sources varied from 3 to 32 and the experiments were done for all three types of fractions and RED. The Fig. 13 shows the performance with mixed TCP and UDP sources.

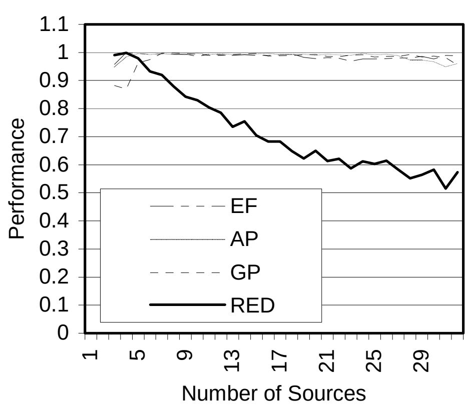

Fig. 12. Performance in Single Congested Link with TCP sources

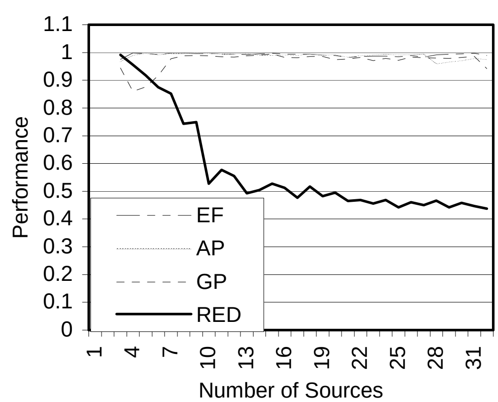

Fig. 13. Performance in Single Congested Link with TCP and UDP sources

In Fig. 13 performance of the three fractions are much better than simple RED. UDP flows are congestion insensitive and TCP flows are congestion sensitive. UDP flows try to get all the bandwidth and so TCP sources get very less bandwidth. Our technique provides good protection of TCP sources from UDP sources and achieves excellent performance. However for RED, TCP flows are not protected and thus RED performs poorly.

4.1 Multiple Congested Links

In the case of multiple congested links, four different sets of experiments were done. The network topology for the simulations is shown in Fig. 14.

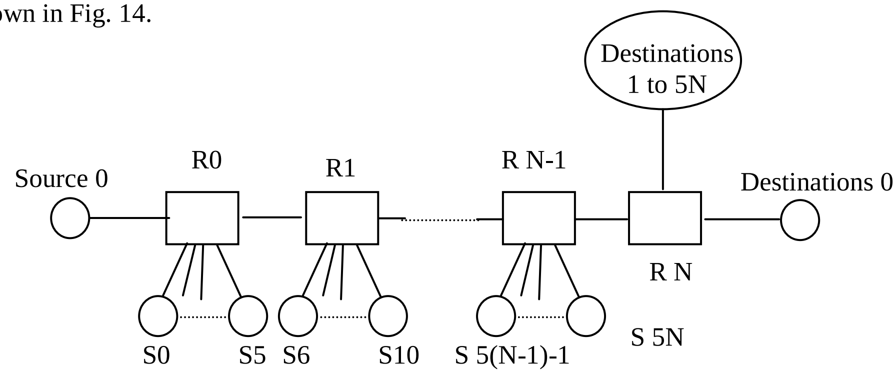

Fig. 14. Multiple Congested Links – Network Configuration

This is a typical parking lot configuration. Flows travel different distances in the network. There are N+1 routers R_0 to R_N. The links connecting the routers have a bandwidth of 10 Mbps and a link delay of 1 ms. At routers R_0 to R_{N-1}, flows enter the network and at the router R_N, all the flows leave the network. At router R_0, source S_0 enters the network. At router R_i sources S_{i*5+1} to $S_{(i+1)*5}$ enter the network. In each experiment set, the number of congested links varied from 2 to 5.
The performance of PAB in multiple congested links is defined as the ratio of the throughput of source S_0 to its SIR divided by the ratio of the sum of the throughputs of all sources to sum of SIRs of all sources. This measure is the allocation ratio of source S_0.

$$AR(0) = \frac{\left[Throughput(0) / SIR(0)\right]}{\left[\sum Throughputs / \sum SIRs\right]} \tag{15}$$

In the experiment set 1, the source S_0 is a UDP CBR source with its SIR set at 5 Mbps. The sources S_1 to S_{n*5} were UDP CBR sources. The SIRs of the sources were set at 5Mbps. The sources were sending data at their SIR. The Fig. 15 shows the AR(0) vs. the number of congested links for equal fractions, AP fractions, GP fractions and RED.

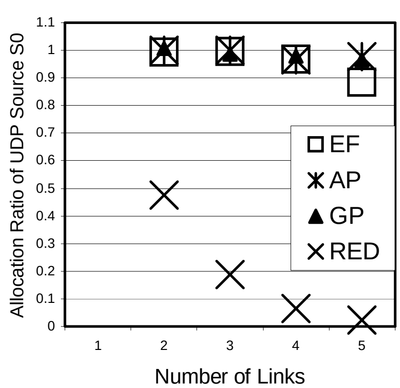

Fig. 15. Allocation Ratio of UDP Source0 in MCL with equal SIR

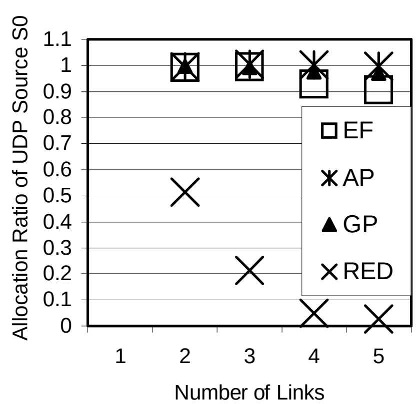

Fig. 16. Allocation Ratio of UDP Source0 in MCL with random SIR

In the experiment set 2, the source S_0 is a UDP source with its SIR set at 5 Mbps. The sources S_1 to S_{n*5} were UDP CBR sources. This experiment is different from experiment set 1, in the nature of the SIR of the sources. The SIR of the sources were set to random values between 1 and 10 Mbps and the sources were sending data at their SIR. The Fig. 16 shows the AR(0) vs. the number of congested links for equal fractions, AP fractions, GP fractions and RED.

From Fig. 15 and Fig. 17, it is clear that the performance of RED is very poor. Among the three types of fractions, EF performs worse than AP fractions or GP fractions. Since during congestion AP fractions and GP fractions are better suited because the high priority labels in their cases are associated with a smaller fraction of the SIRs of the sources. In the experiment set 3, the source S_0 is a TCP Telnet source with its peak rate set at 5 Mbps. All the other sources S_1 to S_{n*5} were UDP CBR sources sending data at their SIR of 5 Mbps similar to experiment set 1. The Fig. 17 shows the AR(0) vs. the number of congested links for equal fractions, AP fractions, GP fractions and RED.

In the experiment set 4, the source S_0 is a TCP Telnet source with its peak rate set at 5 Mbps. All the other sources S_1 to S_{5*n} were the same as in experiment set 2 with their SIRs and data rates varying randomly from 1 to 10 Mbps. The Fig. 18 shows the AR(0) vs. the number of congested links for equal fractions, AP fractions, GP fractions and RED.

The performance of RED with TCP as source S_0 is very poor and almost nil. Among the three types of fractions, performance variation occurs as the number of congested links increases. This is due to the fact that this technique is an approximate implementation of PAB. Further TCP behavior varies widely depending on the threshold value and the actual fraction of the SIR currently allowed through the link.

Further the performance of equal fractions suffers significantly as the number of congested links increases. In equal fractions the highest priority label is associated with a fraction value that is $1/8^{th}$ of SIR. So as severity of congestion increases all packets that are marked with the lower priorities are dropped. Packets of the highest priority alone survive the congestion. Since now all packets are of the same priority, it becomes difficult to achieve PAB.

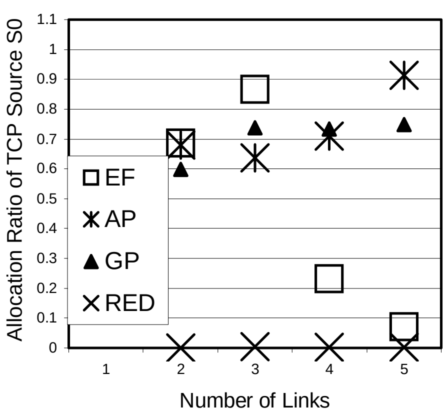

Fig. 17. Allocation Ratio of TCP Source-0 in MCL with equal SIR

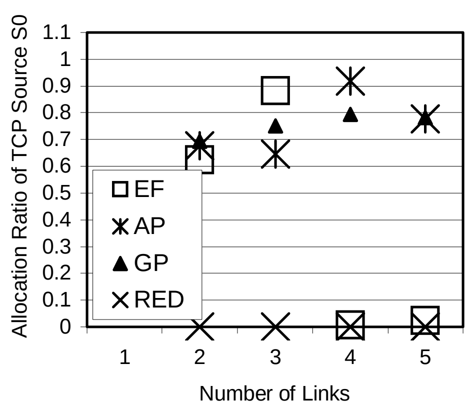

Fig. 18. Allocation Ratio of TCP Source-0 in MCL with random SIR

5. CONCLUSIONS

Proportional allocation of bandwidth is a scheme that takes into consideration the differences in SIR of various flows that are sharing a congested link. In the proportional allocation scheme, bandwidth is allocated in proportion to SIR of the competing flows. Our technique to implement PAB without per–flow state maintenance uses multiple token buckets to label the packets at the edge of the network and multilevel threshold queue at the core routers to discard packets during congestion. The labels are associated with fractions and each label corresponds to a fraction of the SIR of a flow. Our simulations show that the performance of our technique is good in both single congested link and multiple congested links. The fractions can be either all equal, or can form an arithmetic progression or a geometric progression. From our study it is clear that the effectiveness of the three schemes: equal fractions, arithmetic fractions or geometric fractions is highly dependent on the congestion occurring in the network.

ACKNOWLEDGEMENTS

This research was supported in part by a grant from ITEC-Ohio, Columbus, Ohio

REFERENCES

1. I. Stoica, S. Shenker, H. Zhang, “Core-Stateless Fair Queuing : Achieving Approximately Fair Bandwidth Allocations in High Speed Networks”, Computer Communication Review. Vol.28, no.4 Oct. 1998. p. 118-30.
2. A. Demers, S. Keshav and S.Shenker, “Analysis and Simulation of a Fair Queuing Algorithm”, Computer Communication Review SIGCOMM '89 Symposium Communications Architectures and Protocols v 19 n 4 1989 p.1-12.
3. Z. Cao, Z. Wang, E. Zegura, “Rainbow Fair Queuing: Fair Bandwidth Sharing Without Per-Flow State”, Proceedings IEEE INFOCOM 2000,Conference on Computer Communications, p.922-31 vol.2.
4. D. D. Clark, W. Fang, “Explicit allocation of Best-Effort Packet Delivery Service”, IEEE/ACM Transactions on Networking, v 6 n 4 August 1998 p.362-73.
5. S. Floyd and V. Jacobson, “Random Early Detection Gateways for Congestion Avoidance”, IEEE/ACM Transactions on Networking, v 1 n 4 August 1993 p.397-413. August 1993.
6. M. Shreedhar, G. Varghese, “Efficient Fair Queuing using Deficit Round Robin”, IEEE/ACM Transactions on Networking. v 4, n 3. June 1996, p 375-385.
7. S. Floyd and K. Fall, “Promoting the Use of End-to-End Congestion Control in the Internet”, IEEE/ACM Transactions on Networking, v 7, n 4. 1999, p 458-472.
8. S. Floyd and V. Jacobson, “Link Sharing and Resource Management Models for Packet Networks”, v 3, n 4. August 1995, p 365-386.
9. V. Jacobson, “Congestion Avoidance and Control”, Computer Communication Review SIGCOMM '88 Symposium: Communications Architectures and Protocols v 18 n 4 1988 p.314-29.
10. D. Lin, R. Morris, “Dynamics of Random Early Detection”, Computer Communication Review, ACM SIGCOMM 97 Conference. v 27 n 4 1997 p.127-37.
11. B. Suter, T. V. Lakshman, D. Stiliadis, A. Choudhury , “Efficient Active Queue Management for Internet Routers” , Proceedings of INTEROP' 98.
12. B. Braden et al. “Recommendations on Queue Management and Congestion Avoidance in the Internet”, RFC 2309, April 1998.
13. S. Blake et al. “An Architecture for Differentiated Services”, RFC 2475, December 1998.
14. J. Heinanen, R. Guerin, “A Single Rate Three Color Marker”, RFC 2697, September 1999.
15. J. Heinanen, R. Guerin, “A Two Rate Three Color Marker”, RFC 2698, September 1999
16. J. Heinanen, F.Baker, W.Weiss, J.Wroclawski, “Assured Forwarding PHB Group”, RFC 2597. June 1999.
17. M. Goyal, A.Durresi, R.Jain, C.Liu, “Performance Analysis of Assured Forwarding”, IETF Draft.- draft-goyal-diffserv-afstdy-00 - February 2000 - Work in progress.
18. Network Simulator – NS – Available at http://www.isi.edu/nsnam/ns/
19. DiffServ patch for NS, S. Murphy, - Available at http://www.teltec.dcu.ie/~murphys/
20. W. Feng, D. Kandlur, D. Saha and K. Shin, "A Self-Configuring RED Gateway", Proceedings of INFOCOM'99: Conference on Computer Communications 1999, p.1320-8 vol.3.
21. V. Jacobson, "Notes on using RED for Queue Management and Congestion Avoidance", talk at NANOG 13, June 1998. http://www.aciri.org/floyd/red.html
22. T. J. Ott, T. V. Lakshman, and L. H. Wong, "SRED: Stabilized RED", Proceedings of INFOCOM'99: Conference on Computer Communications 1999, p.1346-55 vol.3.
23. A. Charny, D. Clark and R. Jain, Congestion control with explicit rate indication", ICC '95, IEEE International Conference on Communications 1995, p.1954-63 vol.3.
24. S. D. Cnodder, O. Elloumi and K. Pauwels, "Effect of Different Packet Sizes on RED Performance", Proceedings ISCC 2000, Fifth IEEE Symposium on Computers and Communications, 2000, p. 793-9.
25. Z. Wang and J. Crowcroft, "A New Congestion Control Scheme: Slow Start and Search (Tri-S)", ACM Computer Communication Review, vol. 21, p 32-43, Jan 1991.
26. J. C. R. Bennett and H. Zhang, "Hierarchical Packet Fair Queueing Algorithms", IEEE/ACM Transactions on Networking v 5 n 5 1997 p.675-89