

Scalable Proportional Allocation of Bandwidth in IP Satellite Networks¹

A. Durresi, P. K. Jagannathan, R. Jain
The Ohio State University
2015 Neil Avenue
Columbus, OH 43210
614-688-5610
{durresi, pjaganna, jain}@cis.ohio-state.edu

Abstract—Proposed satellite constellation networks, based on Geostationary Earth Orbit (GEO) satellites, Medium Earth Orbit (MEO), and Low Earth Orbit (LEO) systems, will be required to transport IP traffic and to provide quality of service (QoS). This paper proposes a new Diffserv-based scheme of IP bandwidth allocation during congestion, called proportional allocation of bandwidth (PAB). PAB can be used in GEO, MEO and LEO satellite networks. In PAB, during congestion all flows get a share of IP available bandwidth, which is in proportion to their subscribed information rate. We suggest a method for implementing PAB without storing per-flow state, which makes the scheme scalable and simple. We show by simulation the advantages of using PAB in IP satellite networks.

TABLE OF CONTENTS

1. INTRODUCTION	1
2. PROPORTIONAL ALLOCATION OF BANDWIDTH	2
3. IMPLEMENTATION OF PAB	2
4. SIMULATION RESULTS	5
5. CONCLUSIONS	11
REFERENCES	11

1. INTRODUCTION

The rapid globalization of the telecommunications industry and the exponential growth of the Internet are placing severe demands on global telecommunications. This demand is further increased by the convergence of computing and communications and by new applications such as Web surfing, desktop and video conferencing. Satisfying this requirement is one of the greatest challenges before telecommunications industry in the Twenty first century. Satellite communication networks can be an integral part of the newly emerging national and global information infrastructures.

In the past three years, interest in Ka-band satellite systems has dramatically increased, with over 450 satellite applications filed with the ITU [1]. In the U.S., there are currently 13 Geostationary Satellite Orbit (GSO) civilian Ka-band systems licensed by the Federal Communications Commission (FCC), comprising a total of 73 satellites. Two Non-Geostationary Orbit (NGSO) Ka-band systems, comprising another 351 satellites, have also been licensed. Eleven additional GSO, four NGSO, and one hybrid system Ka-band application for license and 16 Q/V-band applications have been filed with FCC [2].

The space segment development has now reached the network layer, and satellites in such constellations will support onboard routing and switching. In this case, the satellite constellation is a true IP network [3, 4]. Along with advantages of using IP protocols, satellite networks will be faced with related problems, with congestions being one of the most serious among them [5, 6]. Satellite systems have also several inherent constraints. The resources of the satellite communication network, especially the satellite and the Earth station, are expensive and typically have low redundancy; these must be robust and be used efficiently. The large delays in Geostationary Earth Orbit (GEO) systems and delay variations in Low Earth Orbit (LEO) systems affect both real-time and non-real-time applications. As a result, QoS issues for broadband satellite networks are somewhat different from those of terrestrial networks.

There has been an increased interest in developing Differentiated Services (DS) architecture for provisioning IP QoS over satellite networks [7]. DS aims to provide scalable service differentiation in the Internet that can be used to permit differentiated pricing of Internet service [8, 9, 10, 11, 12]. In this paper we present a new scheme in which during congestion, IP bandwidth is allocated in proportion to the users' negotiated bandwidth agreement.

¹ 0-7803-7651-X/03/\$17.00 © 2003 IEEE

² IEEEAC paper #1197, Updated December 10, 2002

The current Internet, over terrestrial and satellite networks, provides best effort service. There are used dropping mechanisms, mostly based on the max-min fairness principle, to determine the packets to be dropped during congestion [13]. There are proposed many mechanisms for achieving max-min fairness using per-flow state information. Mechanisms like CSFQ [14, 15, 16, 17], which do not store per-flow state information, have achieved approximate max-min fairness among competing flows.

However in the current Internet (including the satellite part of it), 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 in 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. We define a new method of bandwidth allocation called Proportional Allocation of Bandwidth – PAB, in which bandwidth must be 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 – ingress router in the satellite network, which has information about the source’s SIR. The satellite router during congestion 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.

In terrestrial Internet, over-provisioning is the main solution used against temporary congestions. Such a solution cannot be afforded in satellite networks. PAB’s scalability and simplicity makes it especially appealing for satellite networks, where the overhead and complexity of solutions is less tolerated than in terrestrial networks. So PAB can help to improve the satellite network efficiency as well as the quality of service offered to customers.

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 in different satellite configuration.

The organization of the paper is as follows: Section 2 deals with the operation of PAB-Proportional Allocation of Bandwidth. The Section 3 deals with technique for

implementation of PAB. In Section 4 we discuss about the simulation results for single congested link and multiple congested links for GEO and LEO satellite networks. We conclude in the Section 5.

2. PROPORTIONAL ALLOCATION OF BANDWIDTH

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 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 [14], each flow is allocated bandwidth as given by:

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

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

Here $\text{send}(i)$ 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 in 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:

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

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

Here, $\text{SIR}(i)$ gives 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 $\text{frac} * \text{SIR}$ 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.

3. IMPLEMENTATION OF PAB

Our technique for implementing PAB is based on the principles of Differentiated Services [8, 10]. 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 satellite 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 satellite network, bandwidth allocation is done using these labels to differentiate between packets. The satellite routers perform 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 labels is fixed for all flows, but the number of label values 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 Figure 1. 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 lower priority labels than the flow with the lower FR as shown in Figure 2. 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 Figure 1 and Figure 2, we assume that there are only four labels and all labels are associated with the same fraction value $\frac{1}{4}$.




Figure 1 - Packet Labeling in flows with same FR but different SIR

To label the packets, multiple token buckets are used at the source or ingress router, which has the knowledge about the SIR of the source. 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.

$$\text{Token Rate of Bucket } j = \text{Frac}(j) * \text{SIR} \quad (5)$$

$$\sum \text{Frac}(j) = 1 \quad (6)$$




Figure 2 - Packet Labeling in flows with same SIR but different FR

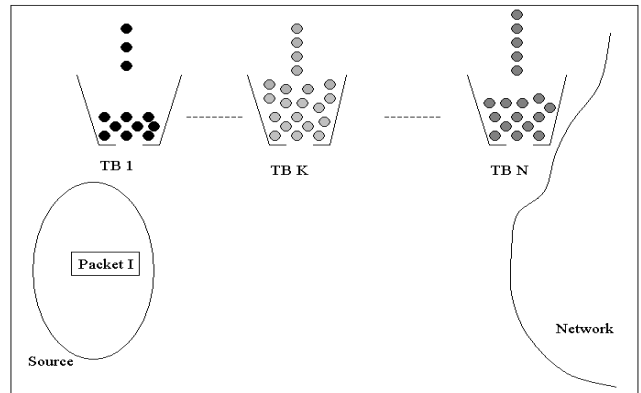


Figure 3 - Before Packet Labeling

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.




Figure 4 - After Packet Labeling

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.

Figure 3 and Figure 4 show a pictorial representation of the token bucket system used for labeling packets. In Figure 3 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 Figure 4. Since packet (I) removes tokens from the token bucket K, packet (I) is marked with label value K.

3.2. Packet Dropping Mechanism at the Core Routers

At the core satellite router, in case of congestion 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 the Multilevel Threshold Based Queuing - MLTQ is similar to SAMT- Single Accounting Multiple Thresholds scheme of [7] and RIO scheme of [13, 18, 19, 20, 21].

In MLTQ, 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 label values. There exist 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 Figure 5. 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.

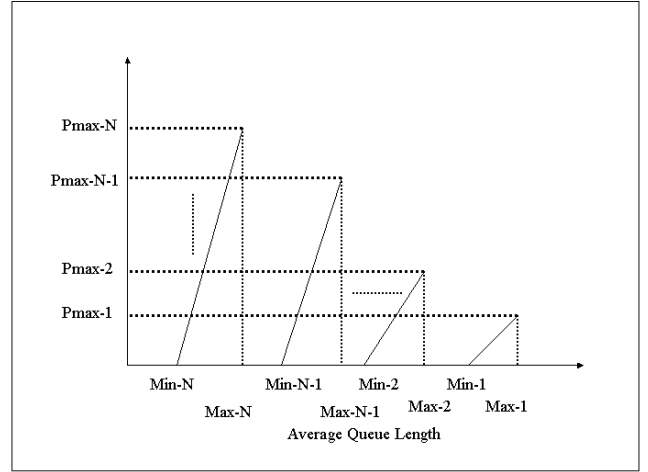


Figure 5 - Core Router Mechanism for Dropping Packets.

3.3. Determination of the label fractions

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

- Fractions with equal value - Equal fractions
- Fractions forming arithmetic progression - AP fractions
- Fractions forming geometric progression - GP fractions

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.

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:

$$a, a+d, a+2d, \dots, a+(N-1)d \quad (7)$$

For simplification, we assume ‘a’ to be equal to ‘d’. So this gives the values of the fractions to be

$$d, 2d, 3d, \dots, Nd \quad (8)$$

Since the sum of the fraction values in the arithmetic progression must be unity, we obtain the following value for ‘d’,

$$d = \frac{2}{N * (N + 1)} \quad (10)$$

For eight labels, the values of fractions are 1/36, 2/36, 3/36, 4/36, 5/36, 6/36, 7/36, and 8/36.

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:

$$a, ar, ar^2, \dots, ar^{(N-1)} \quad (11)$$

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

$$Sum = \frac{a(1 - r^N)}{1 - r} = 1 \quad (12)$$

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:

$$r, r^2, r^3, \dots, r^N \quad (13)$$

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

$$r(2 - r^N) = 1 \quad (14)$$

For N = 8, r gets a value approximately equal to 1/2. 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.

The Figure 6 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 and 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 that of RED. The

performance of PAB in single congested link for GEO satellite networks was studied. Further the performance of PAB in multiple congested links for LEO satellite networks was also studied.




Figure 6. Value of Equal, AP and GP Fractions

4.1. Single Congested Link

We used the ns-2 simulator [22, 23] for performing simulations. The packets are marked with labels at the edge of the network at the ingress routers, which are the gateways to the satellite network. The satellite routers use the labels for providing service and dropping packets, when there is congestion.

In the case of single congested link, the network configuration is as shown in Figure 7. There are N flows sharing a single congested bottleneck link. The bottleneck link is the satellite link between terrestrial router TR1 and the terrestrial router TR2 through the satellite router SR. Same results were obtained when a SR-TR2 link was

congested. The SIR of the i^{th} flow was set at $\left(\frac{i}{N} * 500\right)$

Kbps. The number of flows sharing the link varied from 3 to 32. The capacity of the bottleneck link is 1 Mbps. The link delay for GEO was set to 125 ms one way. The capacity of the link buffer was 100 packets. The packet size of the TCP flows was set at 1000 bytes. The packet size of the CBR flows was set at 210 bytes. The parameters for 8-level drop thresholds at the core router is shown in the Table 1.

For the token buckets at the ingress routers, the bucket size was fixed at 80000 bytes. The token rate for the bucket of every label is the product of the fraction associated with that label and the SIR of the flow.

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.

$$TR(i) = \frac{\text{Throughput of flow}(i)}{\sum \text{Throughput of all flows}} \quad (15)$$

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

$$SR(i) = \frac{SIR \text{ of Source}(i)}{\sum SIR \text{ of Source}(i)} \quad (16)$$

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

$$AR(i) = \frac{TR(i)}{SR(i)} \quad (17)$$

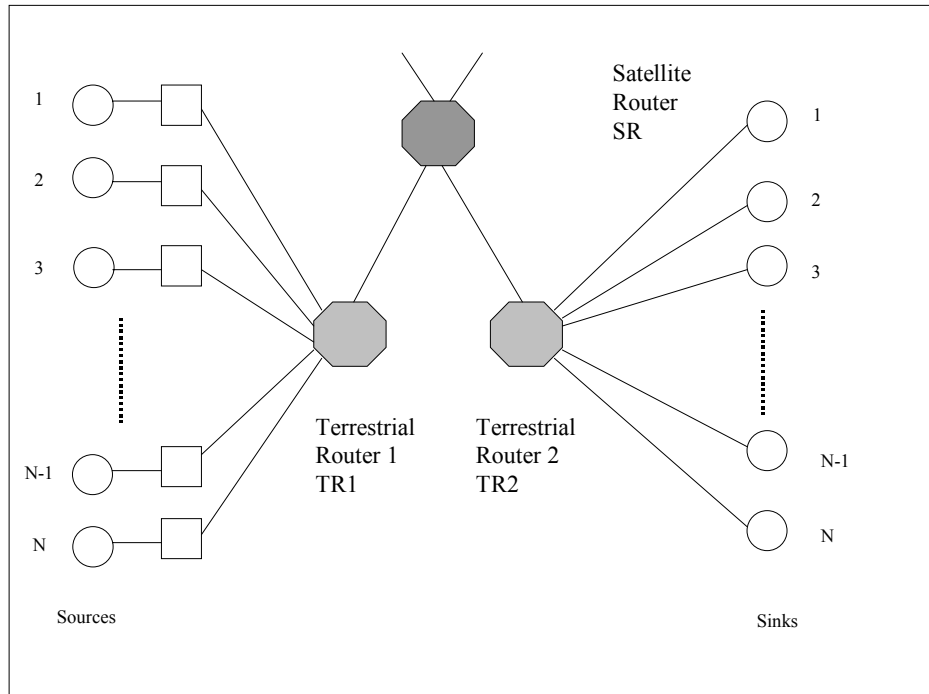


Figure 7 - Single Congested Link in Satellite Network - Network Configuration

Table 1- 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

The performance measure is given by:

$$\text{Proportionality Index} = \frac{[\sum AR(i) * \sum AR(i)]}{N * \sum [AR(i) * AR(i)]} \quad (18)$$

This is used as a measure of fairness in [24]. The proportionality index is inversely related to the difference in the allocation ratios of the various sources. The more the allocation ratios of the various sources differ, the more the value of the proportionality index decreases.

The flows in the network were either UDP or TCP. Experiments were performed with three different combinations of flows. In the first set of experiments all the flows were UDP. Each source sends one flow into the network and the flow reaches the destination at the other end of the network. Experiments were conducted with the number of flows increasing from 3 to 32. The sources were sending data randomly between 10% to 200% of their SIR. The experiments were performed using PAB and RED. Figure 8 shows the results for the experiments. The graph shows the number of flows vs. the proportionality index for UDP flows. The second set of experiments all the flows were TCP. The peak data rate of the TCP flows was set to their SIR. Figure 9 shows the results for the experiments. The graph shows the number of flows vs. the proportionality index for TCP flows.

In the third set of experiments all the even flows were TCP and all the odd flows were UDP. The data rates of the UDP and TCP flows were the same as in the previous experiments. Figure 10 shows the results for the experiments. The graph shows the number of flows vs. the proportionality index for mixed flows. When all flows were UDP it can be observed that as the number of flows increases, PAB performs well. However for RED, the performance drastically drops as the number of flows 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.

In the case of TCP flows, as the number of flows increases the performance of PAB is 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 flows. During congestion the TCP flows reduce the sending data rate so that the rate of all flows are equal and RED cannot distinguish between flows and thus has poor performance.

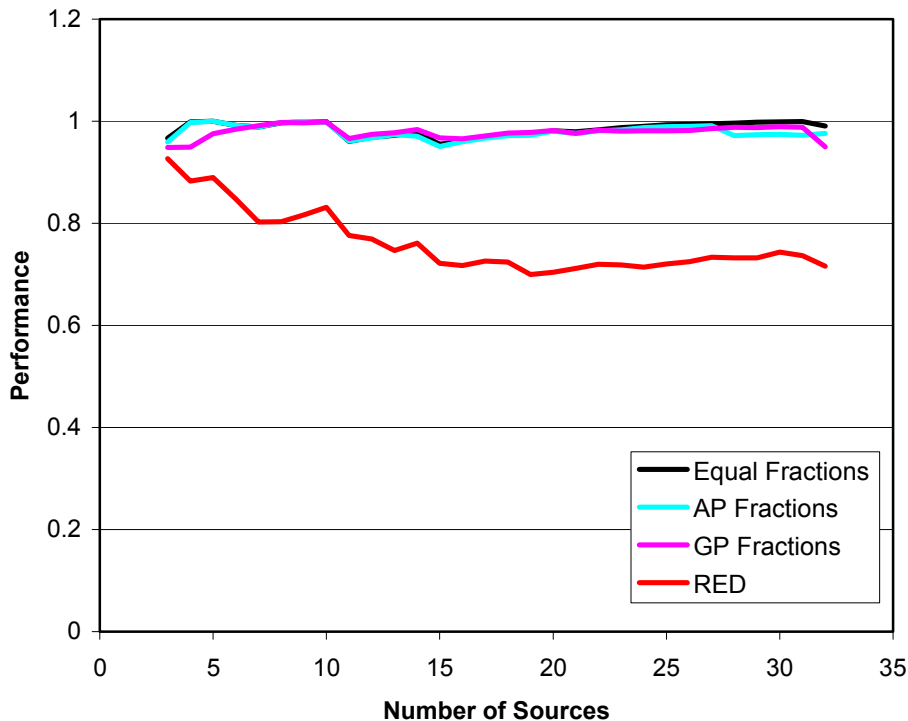


Figure 8- Performance in SCL with UDP flows in GEO satellite networks

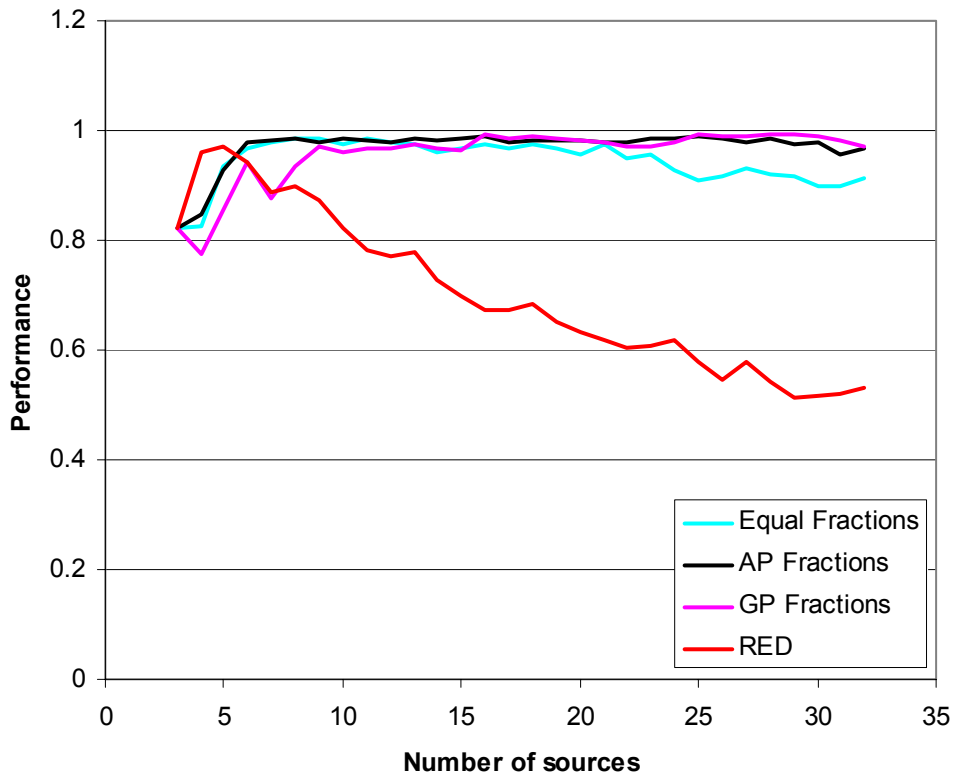


Figure 9-Performance in SCL with TCP flows in GEO satellite networks

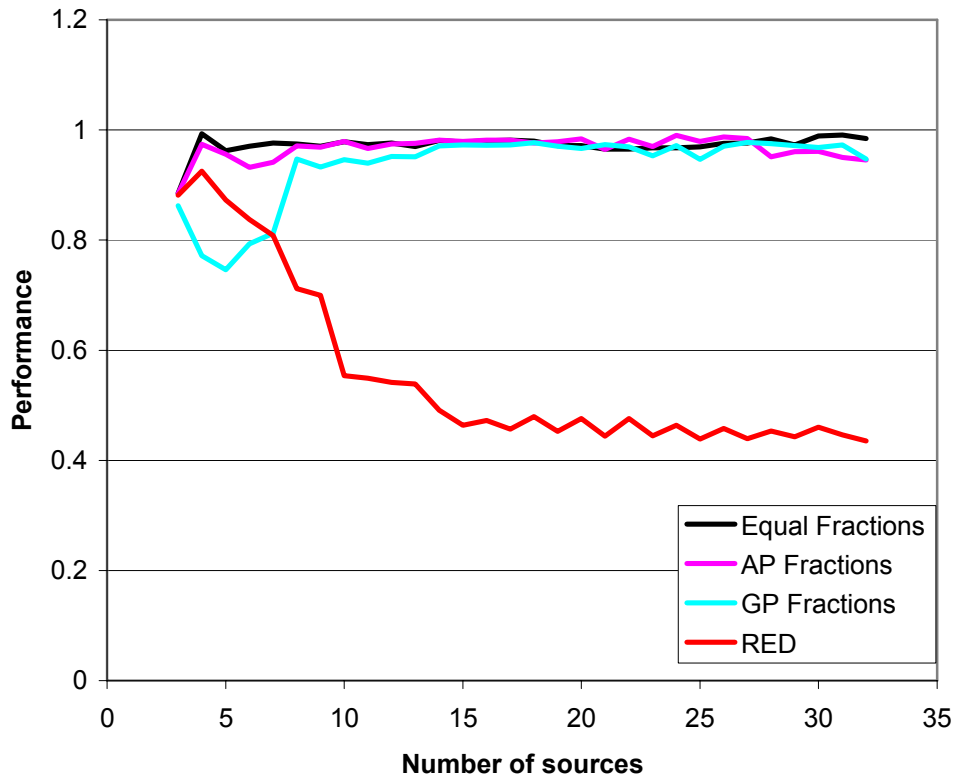


Figure10- Performance in SCL with TCP and UDP flows for GEO satellite networks

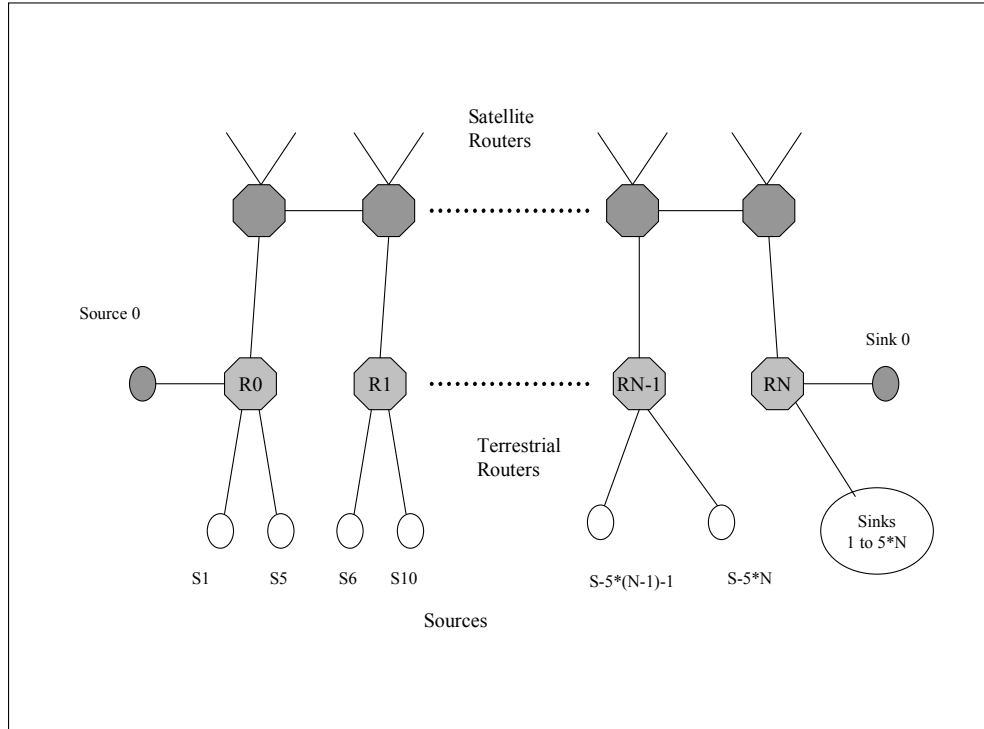


Figure 11-Multiple Congested Links – Network Configuration

For mixed flows, the even flows were TCP and the odd flows were UDP. Both the TCP flows and UDP flows had similar data rate as in the previous experiments. UDP is congestion insensitive and TCP is congestion sensitive. So UDP flows try to get all the bandwidth and therefore TCP flows get very less bandwidth. Our technique provides good protection of TCP flows from UDP flows and achieves excellent performance. However in RED, TCP flows are not protected and thus RED performs poorly.

4.2. Multiple Congested Links

In the case of multiple congested links, the network topology for the simulations is shown in Figure 11.

This is a typical parking lot configuration. Flows travel different distances in the network. There are $N+1$ terrestrial routers. A terrestrial router is connected to the next terrestrial router through the satellite router. The satellite links connecting the routers have a bandwidth of 1.5 Mbps. The link delay for LEO satellite networks was set 25 ms one way. At the 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 , flow S_0 enters the network. At router R_i flows 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. Similar results were obtained using MEO constellations.

The performance of PAB in multiple congested links is defined as the ratio of the throughput of flow S_0 to its SIR divided by the ratio of the sum of the throughputs of all flows to sum of SIRs of all flows. This measure is the allocation ratio of flow S_0 .

$$\text{Allocation - Ratio - } AR(0) = \frac{[\text{Throughput}(0)/\text{SIR}(0)]}{\left[\frac{\sum \text{Throughputs}}{\sum \text{SIRs}} \right]} \quad (19)$$

Two experiment sets were using TCP and UDP as flow S_0 . From Figure 12 it is clear that the performance of RED is poor with S_0 as a UDP flow. From Figure 13 we can observe that the performance of RED with TCP as flow S_0 is very poor and almost nil. For PAB performance variation occurs as the number of congested links increases. This is due to the fact that our technique is only an approximate implementation of proportional allocation of bandwidth by using a limited number of label priorities. Further TCP behavior varies widely depending on the threshold value and the actual fraction of the SIR currently allowed through the link. In case of high level of congestion GP fractions work better because of their higher granularity for high priority labels.

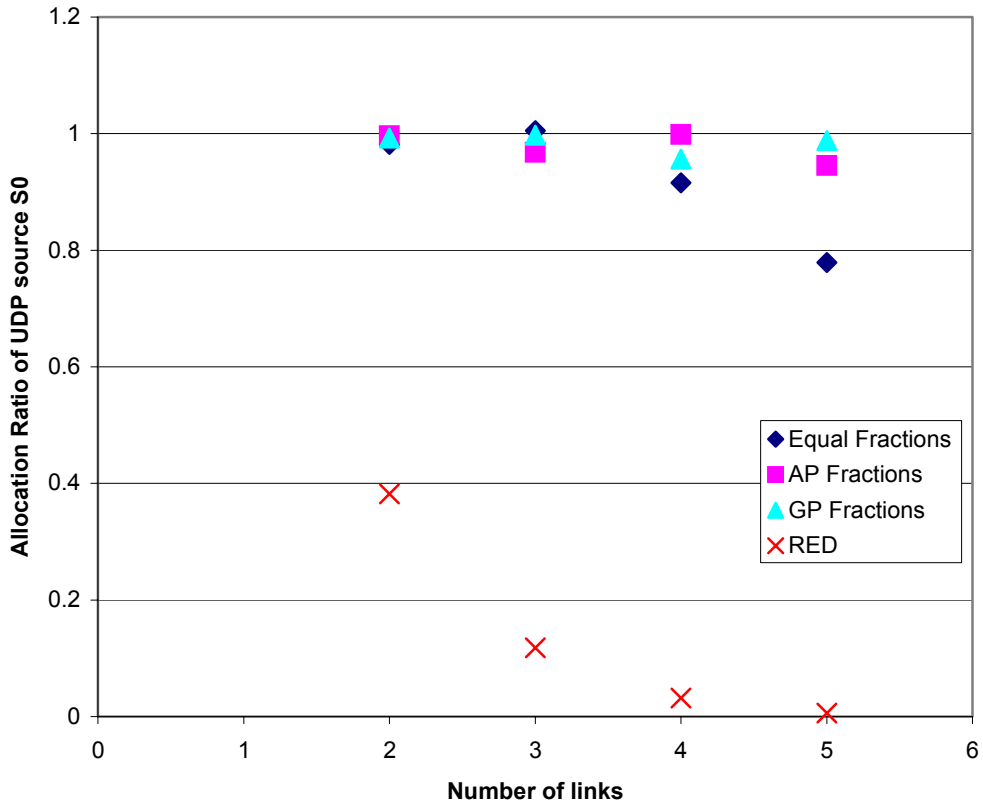


Figure 12 - Allocation Ratio of UDP Flow-0 in MCL for LEO satellite networks

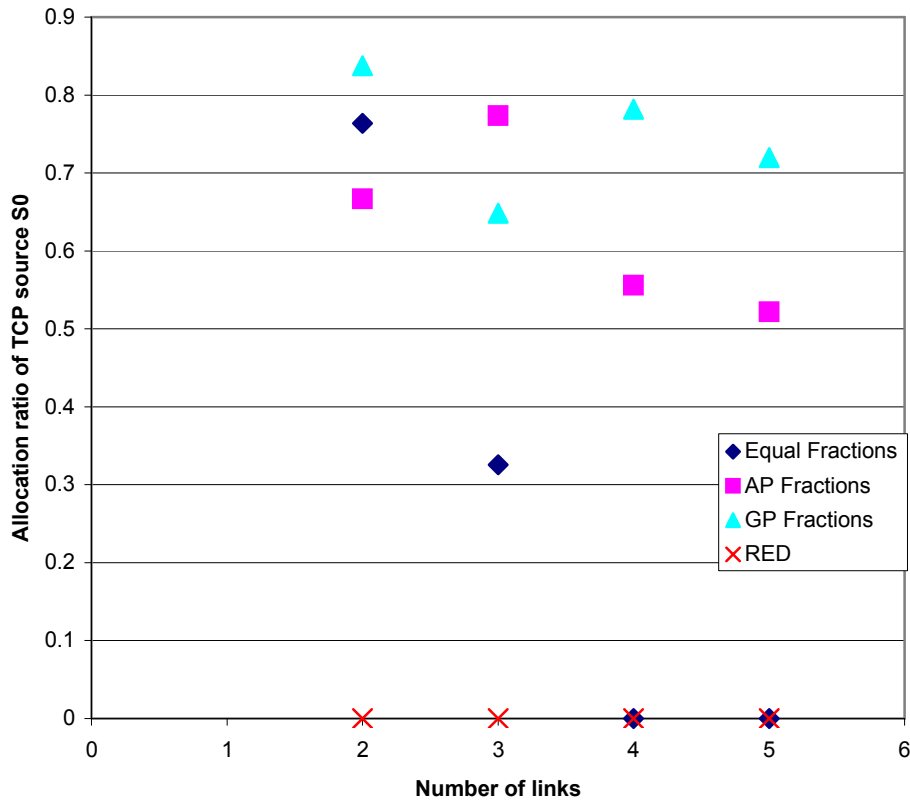


Figure 13 - Allocation Ratio of TCP Flow-0 in MCL for LEO satellite networks

5. CONCLUSIONS

One of the goals of deploying diffserv architecture on a satellite network, either in GEO, MEO or LEO architecture, is to provide scalable service differentiation that can be used to permit differentiated pricing of Internet service.

In this paper we propose a new Diffserv-based scheme of IP bandwidth allocation during congestion, called proportional allocation of bandwidth (PAB). PAB can be used in GEO, MEO and LEO satellite networks. In PAB scheme, bandwidth is allocated in proportion to SIR of the competing flows. We implement PAB, without per-flow state maintenance using multiple token buckets to label the packets at the edge of the network and multilevel threshold queue at the satellite 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 PAB scheme is good in both single congested link and multiple congested links in satellite networks.

REFERENCES

- [1] Sastri Kota, Arjan Duresi, Raj Jain, "Realizing Future Broadband Satellite Network Services," Invited paper in the Symposium *On Modelling and Simulation Environment for Terrestrial and Satellite Networks*, held in Aachen, Germany, 13 – 14 September 2000.
- [2] Sastri Kota, "Multimedia Satellite Networks: Issues and Challenges," Proc. *SPIE International Symposium on Voice, Video, and Data Communications*, Boston, Nov 1-5, 1998.
- [3] L. Wood, A. Clerget, I. Andrikopoulos, G. Pavlou and W. Dabbous, "IP Routing Issues in Satellite Constellation Networks," In *International Journal of Satellite Communications*, Special issue on Internet Protocols over Satellite, vol.18, no.6, Nov/Dec 2000.
- [4] L. Wood, G. Pavlou and B. Evans, "Effects on TCP of Routing Strategies in Satellite Constellations," *IEEE Communications Magazine*, p.172-181, March 2001.
- [5] V. Jacobson, "Congestion Avoidance and Control", *Computer Communication Review SIGCOMM '88 Symposium: Communications Architectures and Protocols* vol. 18, no. 4, 1988 p.314-29.
- [6] S. Floyd and K. Fall, "Promoting the Use of End-to-End Congestion Control in the Internet", *IEEE/ACM Transactions on Networking*, vol. 7, no. 4. 1999, p. 458-472.
- [7] Arjan Duresi, Sastri Kota, Raj Jain, Mukul Goyal, "Achieving QoS for TCP Traffic in Satellite Networks with Differentiated Services", *Space Communications Journal*, vol. 17, no. 1-3, p. 125-136, 2001.
- [8] S. Blake et al. "An Architecture for Differentiated Services", RFC 2475, December 1998.
- [9] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski, "Assured Forwarding PHB Group", RFC 2597. June 1999.
- [10] 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.
- [11] J. Heinanen, R. Guerin, "A Single Rate Three Color Marker", RFC 2697, September 1999.
- [12] J. Heinanen, R. Guerin, "A Two Rate Three Color Marker", RFC 2698, September 1999
- [13] 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.
- [14] 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, p. 118-30, October 1998.
- [15] 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*, vol. 19 no. 4 1989 p.1-12.
- [16] M. Shreedhar, G. Varghese, "Efficient Fair Queuing using Deficit Round Robin", *IEEE/ACM Transactions on Networking*. vol. 4, no. 3. June 1996, p 375-385.
- [17] S. Floyd and V. Jacobson, "Link Sharing and Resource Management Models for Packet Networks", vol. 3, no. 4. August 1995, p 365-386.
- [18] 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.
- [19] D. Lin, R. Morris, "Dynamics of Random Early Detection", *Computer Communication Review*, ACM SIGCOMM 97 Conference. v 27 n 4 1997 p.127-37.
- [20] B. Suter, T. V. Lakshman, D. Stiliadis, A. Choudhury, "Efficient Active Queue Management for Internet Routers", *Proceedings of INTEROP' 98*.
- [21] B. Braden et al. "Recommendations on Queue Management and Congestion Avoidance in the Internet", RFC 2309, April 1998.

[22] Network Simulator – NS – Available at <http://www.isi.edu/nsnam/ns/>

Wesley and “The Art of Computer Systems Performance Analysis” published by Wiley.

[23] DiffServ patch for NS, S. Murphy, - Available at <http://www.teltec.dcu.ie/~murphys/>

[24] R. Jain, “The Art of Computer Systems Performance Analysis”, John Wiley and Sons Inc., 1991.

[25] W.R.Stevens: “TCP/IP Illustrated, The Protocols”. Addison-Wesley, Reading, MA, 1994.



Arjan Duresi is a research scientist at the Computer and Information Science Department and adjunct assistant professor at the Electrical Engineering Department of the Ohio State University. He conducts research in QoS over IP networks, wireless and satellite networks, congestion control, traffic management, optical networking and network security. He has 15 years of experience in industry and academic research. He has a BSEE, MS and Ph.D. from Polytechnic University of Tirana and a Superior Specialization in telecommunications from La Sapienza University in Rome, Italy and Italian Telecommunications Institute.

Prassana K. Jagannathan works at Citrix Systems. He has a BE in Computer Science and Engineering from Anna University, India and a M.S. the Ohio State University



Raj Jain is a professor at the Computer and Information Science Department of the Ohio State University. He is a fellow of IEEE and ACM and serves on the editorial boards of Computer Networks, Computer Communications and the Journal of High Speed Networks. He is the author of two popular books: FDDI Handbook: High Speed Networking using Fiber and Other Media” published by Addison-