Achieving Relative Service Differentiation for TCP Traffic in Satellite IP Networks Using Congestion Pricing

Wei Koong Chai, Kin-Hon Ho, and George Pavlou

Centre for Communication Systems Research, University of Surrey, GU2 7XH, UK
{W.Chai, K.Ho, G.Pavlou}@eim.surrey.ac.uk

Abstract. This paper proposes the use of congestion pricing for achieving relative service differentiation in satellite IP networks. The idea is to use congestion pricing as an effective approach to control traffic rate for Transmission Control Protocol (TCP) flows, with Explicit Congestion Notification as a congestion feedback mechanism, by taking into account network users' willingness-to-pay. With multitude of competing TCP flows, congestion pricing ensures that the higher the user's willingness-to-pay, the higher the traffic throughput. We implement the congestion pricing approach on the *ns-2* simulator and evaluate its performance on both geostationary and non-geostationary satellite IP networks. Simulation results show that congestion pricing can be adopted as an effective approach for service differentiation in satellite IP networks and achieves fair relative service differentiation.

1 Introduction

One of the key objectives of the next generation Internet is to achieve service differentiation. The latter categorizes traffic into a set of classes to provide different treatment in Quality of Service (QoS). In recent years, two well-known architectures, Integrated Services (*IntServ*) [1] and Differentiated Services (*DiffServ*) [2], have been proposed to support service differentiation in the Internet. Meanwhile, satellite networks with vast coverage capability are vital for the future network infrastructure. As satellite networks are carrying increasing volumes of Internet traffic, the ability to provide flexible yet efficient service differentiation will be extremely important.

This paper proposes an approach to achieve *Relative Service Differentiation* (RSD) for satellite IP networks. The RSD model ensures that high-priority traffic will receive no worse QoS than low-priority traffic but it does not guarantee an absolute level of QoS to any traffic class. The key advantage of RSD is its ability to keep service differentiation consistent regardless of network load changes. To realize service differentiation, the authors in [3] claimed that pricing is an indispensable factor to be considered. They argued that proposals that define technologies in isolation of economics and implementation are fundamentally flawed. The authors in [4] also stated that the RSD model must be strongly coupled with a pricing or policy-based scheme to make higher classes more costly than lower ones. However, to the best of our knowledge, RSD that takes into account such an economic incentive for satellite networks has not been investigated. For that reason, we propose to use *congestion pricing* in order to achieve RSD for satellite IP networks.

Y. Koucheryavy, J. Harju, and V.B. Iversen (Eds.): NEW2AN 2006, LNCS 4003, pp. 246 – 258, 2006.

Congestion pricing has been primarily investigated in terrestrial networks as an effective means for resource management and congestion control [5][6]. The fundamental concept of congestion pricing is to inform network users the congestion cost their traffic is incurring via marking and feedback mechanisms such as Random Explicit Detection (RED) [7] and Explicit Congestion Notification (ECN) [8]. The network users then react to these feedback congestion costs by adjusting the traffic sending rate according to their incentive to pay for better QoS. Such incentive is called *Willingness-To-Pay* (WTP). In theory, a user with higher WTP receives better QoS. As such, this becomes a natural mechanism for supporting RSD. Using congestion price for RSD has also an advantage of scalability as it places rate control intelligence at network edges while keeping the core network simple: each network user only reacts to feedback signals by adjusting the traffic sending rate. Although this idea is similar to that proposed for *DiffServ*, the former is much simpler because no complexity is added to the network core routers.

In this paper, we review the state-of-the art of service differentiation in satellite networks and present theoretical background on congestion pricing (Section 2 and 3). We implement congestion pricing as a window-based congestion control, taking into account network users' WTP as a weight to differentiate among their TCP flows in terms of traffic sending rate and throughputs (Section 4). We evaluate our implementation using the network simulator *ns-2* under Geostationary Earth Orbit (GEO) and non-GEO satellite IP network scenarios (Section 5). Simulation results demonstrate that congestion pricing achieves RSD in satellite IP networks through relative and fair bandwidth sharing. Finally, we conclude the paper (Section 6).

2 Service Differentiation in Satellite Networks

A number of proposals have been made for achieving service differentiation in satellite networks. A gateway architecture for IP satellite networks to achieve *DiffServ* via a joint resource management and marking approach is proposed in [9]. Their objectives are to minimize bandwidth wastage while satisfying QoS requirements. In [10], the authors compare several buffer management policies for satellite onboard switching to differentiate real time and non-real time traffic. The feasibility of using multi-protocol label switching for service differentiation is investigated in [11].

In the context of service differentiation for TCP flows, [12] demonstrates this possibility via joint configuration of transport-level and Medium Access Control (MAC)-level service differentiation mechanisms with extensive simulation and analysis. In [13], the authors attempt to realize *DiffServ* for TCP connections with a full-fledged ATM switch onboard with buffer management capacity. The work more related to ours is [14], where an onboard satellite waiting time priority scheduler is proposed to realize proportional differentiated services [15] in a Bandwidth-on-Demand satellite network. To the best of our knowledge, no previous work has considered an economic-incentive approach to achieve relative service differentiation in satellite IP networks, which is the focus of this paper.

3 Theoretical Background of Congestion Pricing

In this section, we briefly review the concept of congestion pricing. For a comprehensive introduction to congestion pricing, readers are referred to [6][16].

We follow the argument in [17] that providing differential QoS implies the use of differential pricing which in turn points to the use of congestion signals to reflect the cost of congestion. Congestion price is dynamic relative to resource loading. When a resource is not saturated, the congestion price is zero. When congestion occurs, the congestion price is an increasing function of extra load on the resource that causes undesirable effects in the network such as increase in congestion and packet loss. Users then react to the congestion price information by adjusting the offered rate according to their WTP.

A formulation of a congestion pricing model is as follows. Suppose that a network consists of a set of nodes and links. The nodes correspond to traffic sources and sinks, while router and link bandwidth are network resources. We denote by J the set of resources and C_j the finite capacity of each resource $j\in J$. Let a route r be a non-empty subset of J between a pair of source and sink, and denote by R the set of all possible routes. We assume that a route is associated with a user and, thus, route and user are used interchangeably. We define a 0-1 matrix $A = (A_{jr}, j\in J, r\in R)$ and set

$$A_{jr} = 1 \quad \text{if route } r \text{ consumes resource } j$$
$$A_{jr} = 0 \quad \text{otherwise}$$

Consider user r has a single flow with his WTP (w_r). Assume that r ranges over the set R and j ranges over the set J, the rate adjustment algorithm [6] (also called *rate-based control*) reacted by users is given by

$$\frac{d}{dt}x_r(t) = \kappa_r\left[w_r - x_r(t)\sum_{j\in r}A_{jr}\mu_j(t)\right]. \tag{1}$$

where

$$\mu_j(t) = p_j\left(\sum_{r\in R}A_{jr}x_r(t)\right) \tag{2}$$

x_r is the offered rate and κ_r is the feedback gain that changes how the rate adjustment behaves. As it can be seen in Eq. (1), the smaller the κ_r the smaller the oscillation on the offered rate. The rate adjustment algorithm is to let users react to the charge of using resources by either increasing or decreasing the offered rate based on their WTP. Suppose that μ_j is the shadow price on using resource j, then the summation term in Eq. (1) represents the path shadow price of the user per unit flow and the term $x_r\sum_{j\in r}A_{jr}\mu_j$ represents the total charge to a user with respect to the offered rate x_r. We follow [5] to interpret $p_j(y_j)$ of Eq. (2) as shadow price of resource j under the load y_j:

$$p_j(y_j) = \frac{d}{dy_j}\delta_j(y_j) \tag{3}$$

where $\delta_j(y_j)$ is the cost function under the load of y_j.

The subtraction of the two terms inside the square bracket in Eq. (1) gives the surplus or loss. The degree of rate adjustment depends on how large is this surplus or loss. If the WTP is higher than the total charge, the offered rate increases; otherwise it decreases. This enables the resource to send feedback signals as a shadow price at rate μ_j to each user who has traffic on it. By adjusting the offered rate, the network attempts to equalize WTP with the total charge at the equilibrium that gives an optimal offered rate. Finally, after each update interval, the offered rate becomes

$$x_r(t+1) = x_r(t) + \frac{d}{dt}x_r(t) \ . \tag{4}$$

In [16], it has been shown that the optimization objective of the above system is to maximize the difference between the total utility over all users (the first term in Eq. (5)) and the total cost incurred on all resources (the second term)

$$\text{Maximize} \quad \sum_{r\in R} w_r \log \lambda_r - \sum_{j\in J} p_j(y_j) \tag{5}$$

where $\lambda_r = \sum_{j\in r} \mu_j$.

4 Window-Based WTP Congestion Control in Satellite Networks

We consider a satellite access network where the bottleneck links are typically satellite up/downlinks. User nodes, associated with different WTP, connect to the satellite network through terrestrial-satellite gateway routers under the control of the satellite operator. As gateways in satellite networks are typically prone to congestion, we employ Random Early Detection (RED) queue as packet marking mechanism and ECN as feedback mechanism. We also use standard drop-tail queues for other nodes. The combined use of RED and ECN enables us to probabilistically mark packets that are causing congestion instead of dropping them.

Based on the rate-based control explained in the previous section, we implement a window-based congestion control algorithm introduced by Gibbens and Kelly [5]. Since TCP is the current dominating protocol in the Internet and its congestion control is window-based, exploring a window-based congestion control algorithm is rational as it avoids drastic changes to the entire Internet framework.

Given that users react to feedback signals by adjusting their offered rate, the TCP congestion window (*cwnd*) can be taken as the effective metric to measure the traffic rate allowed to be sent to the network. Since changing the size of *cwnd* effectively varies the number of packets that is allowed to be transmitted in a Round-Trip Time (RTT), window-based control is closely related to rate-based control [18]. The rate calculated at the source is converted to window size by multiplying it with RTT, i.e. $cwnd_r = x_r RTT$ for each user. Therefore, the rate-based and window-based control mechanisms are compatible to each other.

Initially, *cwnd* is assigned to one segment (or packet) and is adjusted for each received packet acknowledgement (ACK). The transmission of packets in *cwnd* and the receipt of their ACKs are done within a RTT. Hence, we consider a window-based congestion control algorithm [5] derived from Eq. (1) and it adjusts *cwnd* by

$$cwnd_r^+ = \overline{k}_r \left(\frac{\overline{w_r}}{cwnd_r} - f \right) \tag{6}$$

per reception of ACK, where $\overline{k_r} = k_r RTT_r$ and $\overline{w_r} = w_r RTT_r$ denotes the feedback gain and WTP per RTT and $f = \lambda_r$ is the shadow price either equals to 0 if the packet is not marked or equals to 1 if the packet is marked.

Since Eq. (6) takes into account feedback delay with queuing delay as a small component in the RTT and single resource congestion, this fits in nicely to satellite network scenarios where the propagation delay is always the dominant component in the RTT while the satellite links are the congested resources. Hence, we base our congestion control mechanism on Eq. (6). By updating *cwnd*, the congestion pricing approach differs from the conventional TCP congestion avoidance algorithm, which follows double multiplicative decrease i.e. the rate of ECN mark is proportional to the offered rate and *cwnd* is halved when congestion occurs [19].

Our implementation uses ACK packets as congestion indicators to update the *cwnd*. An ECN feedback signal is sent back through ACK to the source to react. We use the four bits ECN in [8] for this purpose. Table 1 details the four bits.

Table 1. ECN bit descriptions

Bit	Description
ECN-capable (ECT)	Set when traffic flow is ECN aware
Congestion Experienced (CE)	Set by the router when it detects an onset of congestion
ECN-echo (ECNecho)	Set by sink if CE bit is set in the received packet header
Congestion Window Reduced (CWR)	Set by source after adjusting its *cwnd* as a response to the ECNecho set

When there is no congestion, the source increases the *cwnd* according to the user's WTP based on Eq. (6) with $f = 0$ (i.e. when no ECN marks received, the congestion window is increased by $\overline{\kappa}_r \overline{w}_r$ every RTT or $(\overline{\kappa}_r \overline{w}_r)/cwnd_r$ for every ACK received since the size of *cwnd* has the equal number of ACKs that can be received within one RTT). During the congestion period, the router marks the packets that cause congestion by setting their CE bit to "1". When the receiver (sink) detects the mark, it sets the ECNecho bit in the ACK packet that it sends back to the source. Upon reception, the source, when knowing that the network resource is getting congested, adjusts its *cwnd* based on Eq. (6) with $f = 1$. The source also sets the CWR bit to indicate that it has responded to the congestion.

The total charge to a user is interpreted as the number of marked packets received per RTT. For instance, if ten packets are sent in an RTT and five of their ACKs are marked, then these marks become the charge per RTT. The condition of marking packets is based on two thresholds of RED: *minthresh* and *maxthresh*. Packets are probabilistically marked if the average queue size is between *minthresh* and *maxthresh* whereby the marking probability is a function of the average queue size. If the average queue size exceeds *maxthresh*, all arriving packets are marked.

5 Performance Evaluation

We implement the window-based congestion control algorithm on the *ns-2* [20]. Our simulation scenarios include GEO and non-GEO satellite networks; shown in Fig. 1 and Fig. 2 respectively. In this paper, we use Low Earth Orbit satellites for non-GEO satellite scenario. Table 2 shows the network parameters used in our simulation.

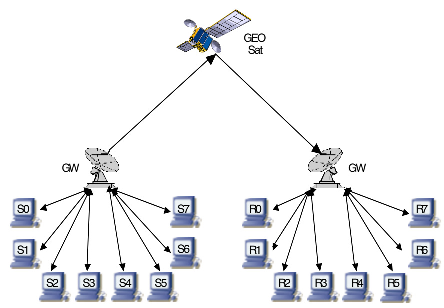

Fig. 1. GEO satellite topology

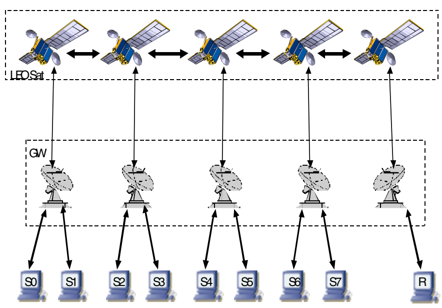

Fig. 2. Non-GEO satellite topology

Table 2. Network Parameters

Link Type	Bandwidth / Delay / Buffer /Queue Type
Up/downlink	1.5Mbyte/s / 125ms(GEO), 25ms(Non-GEO) / 60 packet / RED
Inter-Satellite Link	25Mbyte/s / 1ms / 60 packet / RED
Terrestrial Link	10Mbyte/s / 5μ / 60 packet / DropTail

We select File Transfer Protocol (FTP) as the source application to provide long-duration TCP connections, which ensures that network steady state can be reached. The FTPs are configured to be transmitting packets of size 1Kbytes with packet interval set as 0.005s. For all simulations, we also set $\overline{k_r}$ as 0.05 to avoid large oscillation on the rate adjustment. We initially evaluate our implementation under the assumption that all network links are error-free and then proceed to show that the implementation also adapts to fading loss in the satellite links. As mentioned, we use RED queues for gateway routers and drop-tail queues for other nodes.

The following metrics are used to evaluate the performance of our approach:

- ***TCP throughput*.** we denote t_r and w_r as the TCP throughput and WTP of user r respectively. The algorithm controls the *cwnd* (in segments) of the TCP connections, which directly influences the TCP throughput to be achieved. For this metric, we capture both the instantaneous *cwnd* and TCP throughput over time.

- ***Fairness.*** in theory, users with higher WTP receive higher TCP throughput. Although the RSD model does not restrict the level of QoS guarantees, the resource allocation should be done in a fair manner with respect to WTP. For fairness evaluation, we capture the *Fairness Index* (*FI*) [21] given by

$$FI_r = \frac{FQ_r}{IFQ_r} \tag{7}$$

where $FQ_r = t_r \Big/ \sum_{p=0}^{n} t_p$ is the fairness quotient and $IFQ_r = w_r \Big/ \sum_{p=0}^{n} w_p$ is the ideal fairness quotient. The ideal fairness result is achieved when $FI_r = 1.0$. If $FI_r < 1.0$, the sources get lower expected traffic rate with respect to their WTP, and vice versa.

5.1 GEO Satellite System Scenario

TCP Throughput. We start eight concurrent FTP sources to saturate the satellite uplink and thus create a congestion scenario in the GEO satellite network. Each source node has an FTP application. To evaluate service differentiation, we assign different WTP values to the users, starting from 1.0 with a step up value of 0.5[1]. We compare our implementation against TCP Reno, one of the major TCP implementations for the current best-effort Internet.

Fig. 3 shows the instantaneous *cwnd* of TCP Reno. As can be seen, TCP Reno does not support service differentiation. It treats all the TCP connections as equal: there is no facility for differentiating among TCP connections. This also applies to the other best-effort based TCP variants. In contrast, as shown in Fig. 4, congestion pricing manages to differentiate the sources according to their WTP. In this case, higher

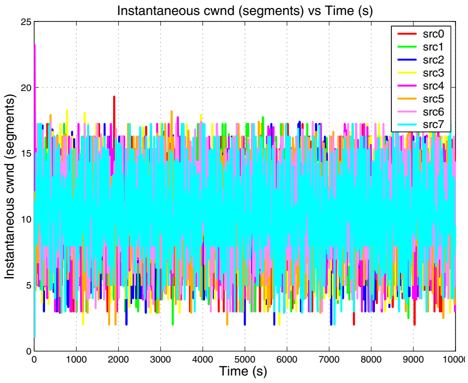

Fig. 3. Standard TCP does not provide differentiation to different flows

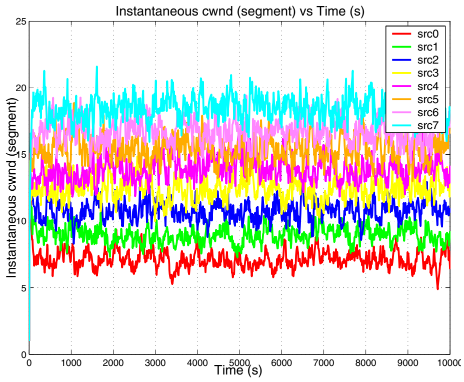

Fig. 4. Congestion pricing manages to differentiate TCP flows

[1] In simulation graphs, src0 has WTP of 1.0, src1 has 1.5, src2 has 2.0, src3 has 2.5 and so on.

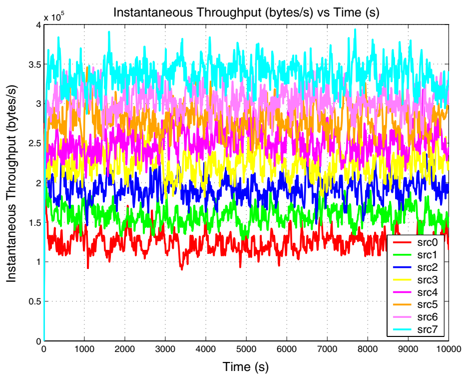

Fig. 5. Differentiated TCP throughput according to user's WTP

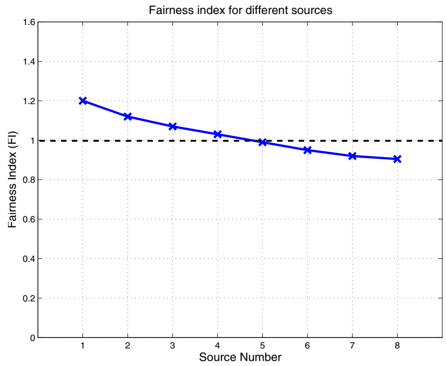

Fig. 6. Fairness index

WTP users consistently have higher instantaneous *cwnd*. With higher instantaneous *cwnd*, higher traffic sending rates are allowed for these sources and, therefore, higher throughput can be achieved. Fig. 5 shows different levels of instantaneous TCP throughput achieved for the FTP sources with different WTP. Fig. 4 and Fig. 5 demonstrate that congestion pricing achieves RSD.

Fairness. We proceed to evaluate fairness achieved by the congestion pricing approach. Fig. 6 shows that the *FI* for most of the sources is close to the ideal case. However, when comparing the results with those produced for terrestrial networks [17], we found that the degree of achieved fairness is lower for GEO satellite systems. This is attributed to the long propagation delay of satellite links, which results in slow feedback of the congestion signal to the sources. The reason that the *FI* for high WTP sources is lower than 1.0 can be explained as follows: as packets are marked continuously because of congestion, higher WTP sources would naturally receive more marked packets in proportion to the number of sent packets. As such, they would receive congestion signals in bulk and eventually lead them to further reduce their *cwnd*. In contrast, low WTP sources receive relatively less congestion feedback. Hence, the gentle slope of the performance for GEO satellite system appears. Together with the TCP throughput results presented in the previous section, congestion pricing achieves fair RSD.

Equal WTP. Previous evaluations are based on the scenario where each user has different WTP. In theory, users with the same WTP should receive identical TCP connection treatment. In this section, we investigate whether congestion pricing achieves equal performance for those users under a congestion scenario. We repeat the previous simulation with a small change. We assign equal WTP of 1.0 to the first two FTP sources. Fig. 7 shows that the two sources having the same WTP receive the same throughput. This ensures that congestion pricing produces reliable and consistent performance in conformance to RSD model.

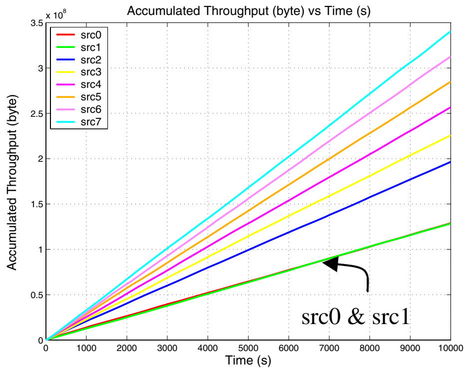

Fig. 7. Sources with same WTP receive same treatment in GEO satellite network scenario

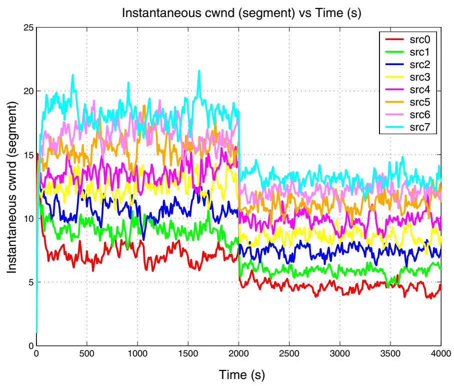

Fig. 8. A loss in the link condition under GEO satellite network scenario

Lossy Satellite Links. We investigate the consistency and adaptability of our approach on satellite links with fading loss. Satellite networks, particularly those using the *Ka band*, are especially susceptible to fade attenuations of signals (e.g. due to rain or bad weather conditions). Two of the more dominant fade countermeasures are *Automatic Repeat Request* (ARQ) and *Forward Error Correction* (FEC). Due to the extra delay incurred by ARQ techniques, its use is basically not recommended [22]. In this paper, we focus on the approach of counteracting fading by the use of adaptive FEC mechanisms which recover erroneous packets that is caused by channel degradation by adding redundancy to packets at the physical layer. The redundancy overhead added is based on the level of signal attenuation measured whereby the redundancy of the packets is increased when C/N (*Carrier/Noise Ratio*) decreases. Therefore, the attenuation effect on a satellite link can be modelled as a decrease of bandwidth [23], as a certain amount of bandwidth has been devoted to carry redundancy overhead rather than information bit. We follow [24] to define the redundancy coefficient for source *i as* $rc_i = IBR_{cs} / IBR_m$ where IBR_{cs} is the information bit rate under clear sky condition while IBR_m is the information bit rate measured at the specific instance. The bandwidth reduction factor is defined as $\phi_i = rc_i^{-1}$. Thus, the actual bandwidth dedicated to the transmission of information bits is given as

$$BW_i^{real} = BW_i \cdot \phi_i ; \qquad \phi_i \in [0,1] \tag{8}$$

where BW_i is the overall bandwidth of source *i*. We run our simulation with all sources set with initial value ϕ_i = 1.0 (clear sky condition) and change the bandwidth reduction factor of all sources to 0.6667 at time, *t* = 2000s to see how the congestion pricing implementation adapts to a change in the link condition. As it can be seen in Fig. 8, at *t* = 2000s, congestion pricing manages to adapt to the changes rapidly (in seconds) while sustaining a consistent and fair RSD. Note that in real world, the fading level may change rapidly, causing the fade countermeasures to continuously adapt to the new fading level. Our simulation represent the case where fading levels are segregated into *fading classes* in which the countermeasure parameter are fixed for all those fading levels within the same class. Hence, our simulation corresponds to the scenario of a transition from one *fading class* to another.

5.2 Non-GEO Satellite System Scenario

We repeat the previous experiments for non-GEO satellite systems. Unlike GEO satellite, the topology of non-GEO satellite is a constellation that connects satellites by inter-satellite links. Note that since our focus is on service differentiation, we ignore the issues arising from satellite mobility such as TCP connection handover.

TCP Throughput and Fairness. Fig. 9 and Fig. 10 show the instantaneous *cwnd* and throughput achieved by each source under a non-GEO satellite network scenario. The service differentiation achieved can be clearly seen. Compared to the GEO satellite scenario (Fig. 4 and Fig. 5), the achieved instantaneous *cwnd* and throughput in the non-GEO satellite network scenario are much lower. This is attributed to the shorter propagation delay in non-GEO satellite systems, which accelerates the congestion feedback from the network to the sources. As such, sources can rapidly react to the feedback signals, reducing *cwnd*.

On the other hand, for the evaluation of fairness, we plot the *FI* achieved by each source in Fig. 11. Compared to GEO satellite systems, it is clear that congestion pricing for non-GEO satellite systems achieves near-ideal fairness as the systems have much shorter propagation delay that enables faster reaction to congestion with more up-to-date congestion feedback signals.

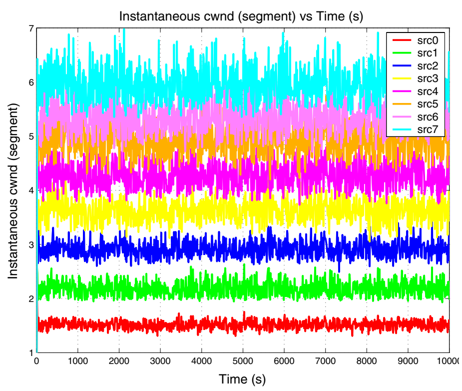

Fig. 9. Differentiated TCP throughput based on user's WTP in non-GEO satellite system

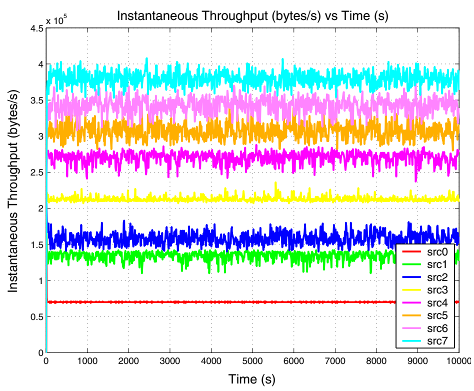

Fig. 10. Differentiated TCP throughput in the non-GEO satellite system

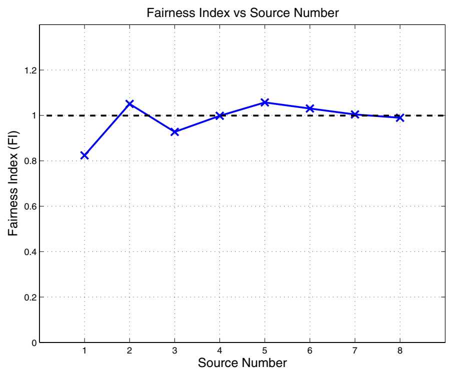

Fig. 11. Congestion pricing achieves fair bandwidth sharing in non-GEO satellite system

Equal WTP and Lossy Satellite Links. Fig. 12 and Fig. 13 show the performance of equal WTP and lossy satellite links under non-GEO satellite network respectively. We found that the results are similar to those under the GEO satellite system scenario. The sources with equal WTP consistently receive approximately equal TCP throughputs. As for lossy satellite link, the congestion pricing approach manages to achieve and maintain RSD among different sources. The plots basically follow the same trend and behaviour as those in the GEO satellite system scenario.

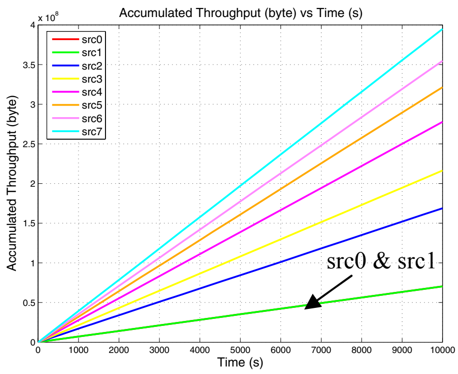

Fig. 12. Sources with same WTP receive same treatment in non-GEO satellite network

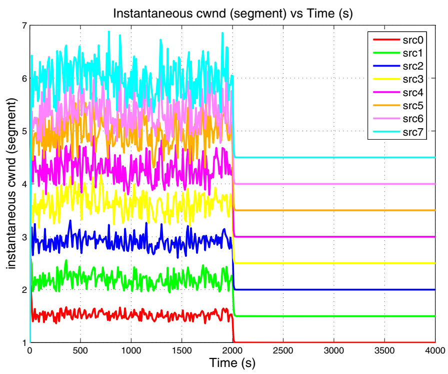

Fig. 13. A loss in the link condition under non-GEO satellite network scenario

6 Conclusions

An RSD scheme using congestion pricing is proposed for satellite IP networks. We have implemented the congestion pricing approach using a weighted window-based congestion control algorithm that takes into account users' WTP to react to congestion costs fed back through an ECN mechanism. Our simulation experiments include both GEO and non-GEO satellite network scenarios. Simulation results show that, for the GEO satellite network, congestion pricing achieves a fair RSD among users with different WTP and this is also the case for a network with lossy satellite links. In addition, we observe that TCP flows with the same WTP receive equal treatment in terms of throughput. Similar conclusions were also drawn for non-GEO satellite networks. However, due to shorter link propagation delay, the achieved performance such as *cwnd* and throughput are smaller than that of the GEO satellite network. In addition, we observe that in non-GEO satellite network, near-ideal fairness can be achieved. As future work, we will investigate service differentiation using Multi-level ECN which allows multiple levels of congestion to be disseminated. This will provide more accurate network status for users to react to their offered rates.

Acknowledgement

This work has been performed within the framework of the Information Society Technologies Research Framework Program 6 (FP6) Network of Excellence SatNEx project, funded by the European Commission (EC). The financial contribution of the EC towards this project is greatly appreciated and acknowledged.

References

1. R. Braden, D. Clark and S. Shenker, “Integrated Services in the Internet architecture: an overview,” *RFC 1633*, June 1994.
2. S. Blake et al., “An architecture for Differentiated Service,” *RFC 2475*, December 1998.
3. P. Key, “Service Differentiation: Congestion Pricing, Brokers and Bandwidth Futures,” Proc. *ACM NOSSDAV’99*, , 1999.
4. C. Dovrolis, and P. Ramanathan, “A case for relative differentiated services and the proportional differentiation model,” *IEEE Network*, September / October 1999, pp. 2-10.
5. R. J. Gibbens and F. P. Kelly, “Resource pricing and the evolution of congestion control,” *Automatica*, 35, 1998.
6. F. P. Kelly et al., “Rate control for communication networks: shadow prices, proportional fairness and stability,” *Journal of the Operation Research Society*, 49, 1998, pp. 237-252.
7. S. Floyd and V. Jacobson, “Early Random Detection Gateways for Congestion Avoidance,” *IEEE/ACM Transaction on Networking*, vol. 1, no. 4, 1993, pp. 397-413.
8. K. K. Ramakrishnan, S. Floyd and D. Black, “The Addition of Explicit Congestion Notification (ECN) to IP,” *RFC 3168*, September 2001.
9. L. S. Ronga et al., “A gateway architecture for IP satellite networks with dynamic resource management and DiffServ QoS provision,” *Int’l Journal of Satellite Communications and Networking*, no. 21, 2003, pp. 351-366.
10. N. Courville, “QoS-oriented traffic management in multimedia satellite systems,” *Int’l Journal of Satellite Communications and Networking*, no. 21, 2003, pp. 367-399.
11. T. Ors and C. Rosenberg, “Providing IP QoS over GEO satellite systems using MPLS,” *Int’l Journal of Satellite Communications and Networking*, no. 19, 2001, pp. 443-461.
12. M. Karaliopoulos et al., “Providing differentiated service to TCP flows over bandwidth on demand geostationary satellite networks,” *IEEE JSAC,* vol. 22, no. 2, 2004, pp. 333-347.
13. A. Duressi et al., “Achieving QoS for TCP traffic in satellite networks with differentiated services,” *Space Communications*, vol. 17, no. 1-3, 2001, pp. 125-136.

14. W. K. Chai et al., "Scheduling for proportional differentiated service provision in geostationary bandwidth on demand satellite networks," Proc. *IEEE GLOBECOM*, 2005.
15. C. Dovrolis et al., "Proportional differentiated services: delay differentiation and packet scheduling," *IEEE/ACM Transactions on Networking,* vol. 10, no. 1, 2002, pp. 12-26.
16. F. P. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, 8, 1997, pp. 33-37.
17. P. B. Key et al., "Congestion pricing for congestion avoidance," Microsoft Research Technical Report MSR-TR-99-15, MSR, 1999.
18. F. P. Kelly. *Mathematical modeling of the Internet "Mathematics Unlimited – 2001 and Beyond"*. Springer-Verlag, Berlin 2001, pp. pp. 685-702.
19. V. A. Siris, C. Courcoubetis and G. Margetis, "Service differentiation and performance of weighted window-based congestion control and packet marking algorithms in ECN networks," *Computer Communications*, vol. 26, 2003, pp. 314-326.
20. Network Simulator-ns (ver. 2) [Online]. Available at http://www.isi.edu/nsnam/ns/
21. M. Shreedhar and G. Varghese, "Efficient Fair Queuing using Deficit Round Robin," Proc. *ACM SIGCOMM*, 1995.
22. M. Marchese and M. Mongelli, "On-line bandwidth control for quality of service mapping over satellite independent service access points," to appear in *Computer Networks*, 2006.
23. R. Bolla, F. Davoli and M. Marchese, "Adaptive bandwidth allocation methods in the satellite environment," Proc. *IEEE ICC*, 2001, pp. 3183-3190.
24. N. Celandroni et al., "An overview of some techniques for cross-layer bandwidth management in multi-service satellite IP networks," Proc. *IEEE GLOBECOM Workshop Advances in Satellite Communications: New Services and Systems*, pp. WO4.4:1-WO4.4:6, 2005.