Providing Relative Service Differentiation to TCP Flows over Split-TCP Geostationary Bandwidth on Demand Satellite Networks

Wei Koong Chai[1], Merkourios Karaliopoulos[2], and George Pavlou[1]

[1] Centre for Communication Systems Research, University of Surrey, GU2 7XH, UK
{W.Chai, G.Pavlou}@surrey.ac.uk
[2] Department of Computer Science, University of North Carolina, Chapel Hill, USA
mkaralio@cs.unc.edu

Abstract. We propose a combined transport – medium access control (MAC) layer scheme to provide *relative* service differentiation to Transmission Control Protocol (TCP) flows over a geostationary (GEO) bandwidth on demand (BoD) satellite networks. Our approach involves the joint configuration of TCP-Performance Enhancing Proxy (TCP-PEP) agents at the transport layer and the scheduling algorithm controlling the resource allocation at the MAC layer. The scheme is independent of the TCP variant used in the network. Extensive simulation results show that the two mechanisms exhibit complementary behavior in achieving the desired differentiation throughout the traffic load space: the TCP-PEP agents control differentiation at low system utilization, whereas the MAC scheduler becomes the dominant differentiation factor at high load.

Keywords: relative service differentiation, satellite network, TCP, Bandwidth on demand.

1 Introduction

Satellite networks are vital components of the next-generation Internet. In order to seamlessly integrate with the global information infrastructure, they have to adjust to the technologies and trends that are adopted in terrestrial networks. For example, even if the DVB/MPEG-2 stack is now the basis for many operational systems, the dominance of Internet Protocol (IP) renders native-IP systems more attractive. All the same, Transmission Control Protocol (TCP) is the Internet de-facto transport protocol for communications and has to be efficiently supported over satellite networks.

Pretty much the same considerations dictate that satellite networks provide service differentiation to different types of traffic in agreement with the Internet quality of service (QoS) framework. The research effort on Internet QoS has been tremendous during the last 15 years and has sparked endless discussions addressing even its necessity as such. More recently, the efforts have been oriented towards frameworks that provide *relative* service differentiation, compromising effectively the hard, quantitative guarantees of Integrated Services (*IntServ*) with the scalability of Differentiated Services (*DiffServ*). Proportional Differentiated Services (PDS) [1] are one of the well-received proposals in this direction.

F. Boavida et al. (Eds.): WWIC 2007, LNCS 4517, pp. 17 – 29, 2007.

How well can satellite networks satisfy the requirements of these lighter QoS frameworks? Which are those functions and mechanisms within the network that can help achieve that? In this paper, we focus on providing *relative* service differentiation to TCP flows over a Geostationary (GEO) Bandwidth on Demand (BoD) satellite network. We assume that the satellite network is equipped with Performance Enhancing Proxies (PEP), which have been widely deployed in satellite networks despite security and scalability concerns [2]. Our approach involves both the transport and medium access control (MAC) layers. TCP flows are divided in a number of classes over the network. To provide them with *relative* service differentiation, we jointly configure the TCP-PEP agents and BoD scheduling algorithm.

The contribution of our work is highly methodological and adds to a broader study of service differentiation mechanisms over GEO satellite networks for all types of traffic [3]. We illustrate the advantages that modest use of cross-layer approaches can have in satellite network engineering. Likewise, we add some arguments in favor of the utility of TCP-PEPs in the same context. However, we equally insist on the evaluation of our proposal via extensive simulations that aim at showing its potential but also its weaknesses.

We organize the paper into six sections. We present the reference system architecture with details on the BoD operation in Section 2. In Section 3, we outline the design requirements for our combined transport–MAC layer differentiation scheme and detail its implementation. We evaluate our proposal in Section 4, showing how the mechanisms at the two layers can be complementary in realizing the desired differentiation objective. Related work is reviewed in Section 5. We conclude the paper in section 6.

2 Network Architecture and BoD Process

2.1 Network Architecture

The system architecture under consideration is a broadband GEO satellite IP network with resource allocation mechanisms analogous to the Digital Video Broadcasting - Return Channel via Satellite (DVB-RCS) standard [4]. We consider satellites with on-board processing allowing mesh terminal connectivity and satellite terminals equipped with TCP-PEPs. Note, however, that our approach is not limited to DVB-RCS like networks. Fig. 1 illustrates the main nodes of the architecture:

- Satellite – the scheduler is assumed onboard.
- Satellite Terminal (ST) – STs are equipped with TCP-PEP agents that split the TCP connections into a terrestrial (wired) component and a satellite (wireless) component. The STs may serve one (residential) or more users (collective).

The multiple access scheme in the satellite link is multi-frequency TDMA (MF-TDMA). The basic unit of link capacity in a MF-TDMA frame is the timeslot (TS), with multiple TSs grouped in TDMA frames along several frequency carriers.

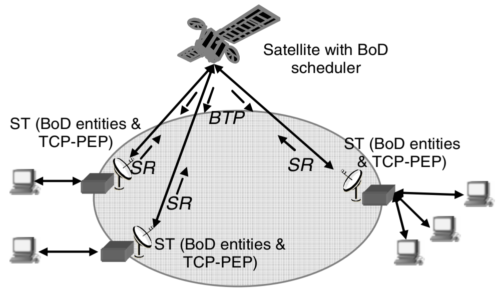

Fig. 1. Reference satellite network configuration

2.2 Bandwidth on Demand (BoD) Process

The BoD process used is drawn from [5] and consists of two functions that are executed periodically, namely the resource request estimation and resource allocation processes. The main entities involved are the BoD entity and the BoD scheduler. The BoD entity is located at the ST and handles all packets of the same class that share the same queue; there will be x BoD entities in a ST supporting x traffic classes. When there are new packet arrivals at their queues, BoD entities send slot requests (SRs) to the BoD scheduler with a period of n_s TDMA frames. If $q(k)$ are the queued packets at the BoD entity at the start of the k^{th} allocation period, then the SR sent to the BoD scheduler is given by

$$SR(k) = \left[\left(q(k) - n_s.a(k) - n_s.\sum_{j=1}^{L-1} SR(k - L_s + j) - n_s.w(k) \right) \Big/ n_s \right]^+ . \tag{1}$$

In Eq. 1, $a(k)$ denotes the number of TS per frame already allocated to the BoD entity for the coming n_s frames. $w(k)$ are the owed TSs by the scheduler from previous resource allocation periods, when the BoD entity requests were partially or not at all satisfied and L_s is the *nominal system response time* (BoD cycle) in frames, namely the time elapsing from the moment a BoD entity submits a request till the moment it should receive its allocation from the scheduler. The L_s parameter accounts for the propagation, transmission and processing delays at the BoD entity and the BoD scheduler. The actual system response time may well be higher than the nominal one, if the request cannot be served in the first encountered resource allocation period but rather has to be queued and served in subsequent resource allocation periods. $[v]^+ = v$ if $v > 0$ and 0 otherwise. This ensures that no SR will be submitted if it is zero or negative.

Upon reception of SRs, the BoD scheduler allocates TSs to each requesting BoD entity based on the scheduling discipline and policies set by the network operator. It then constructs the burst time plan (BTP) that contains the allocation information and broadcasts it to the BoD entities. Fig. 2 outlines the time evolution of the BoD process and the timing of the resource request submission and allocation tasks.

At the STs, TCP-PEPs split the TCP connections between the terrestrial and satellite domains. The proxies will cache TCP segments and prematurely acknowledge their arrival. On the satellite network side, they are required to execute our transport layer differentiation mechanism.

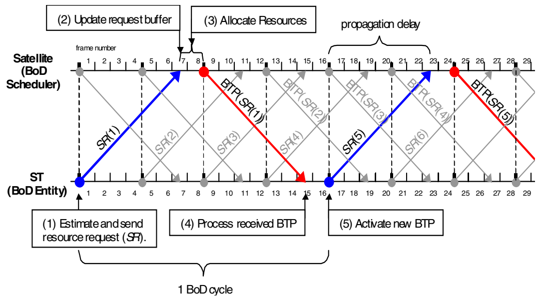

Fig. 2. BoD timing diagram

3 PDS Provision to TCP Traffic over Satellite

3.1 Design Requirements

Our design follows the objective of the PDS model [1]. The model is especially suitable for satellite networks as it is lightweight and does not require complex mechanisms such as admission control or provisioning.

The network considered supports N service classes indexed by i, $i \in I \equiv \{1..N\}$. Each class is assigned a differentiation parameter (DP) r_i controlling the performance gap between service classes. If σ_i denotes the performance metric of interest for class i, then the PDS model requires that

$$\frac{\sigma_i}{\sigma_j} = \frac{r_i}{r_j} \qquad \forall i, j \in \{1, \ldots, N\}. \tag{2}$$

In this paper, class 1 is the highest priority class with its DP set to unity. We normalize all DPs with reference to it so that

$$0 < r_N < r_{N-1} < \ldots < r_2 < r_1 = 1$$

The metric of interest here is the average TCP throughput. We want to provide different throughput to TCP flows that are classified under different service classes. Applying the PDS model into our problem, we would like to have

$$\frac{\overline{thr_i}}{\overline{thr_j}} = \frac{r_i}{r_j} \qquad \forall i, j \in \{1 \ldots N\} \tag{3}$$

where $\overline{thr_i}$ is the average throughput that traffic flows of class i obtain in their lifetime. From here onwards, we refer to the ratio $\overline{thr_i} / \overline{thr_j}$ simply as throughput ratio.

3.2 Combined Transport-MAC Layer Differentiation

To control TCP throughput, we first need to consider what parameters affect it. Existing analytical approximations for TCP throughput under congestion loss, e.g. [6], [7], are most useful to this end

$$thr = \frac{W \cdot MSS}{RTT} = \frac{k}{\sqrt{p}} \cdot \frac{MSS}{RTT} \tag{4}$$

where *W* is the average TCP send window, *MSS* is the maximum segment size of the TCP connection, *p* is the loss probability and *k* is a constant depending on the nature of loss and the use of the Delayed Acknowledgements option [6]. The term *RTT* denotes the overall round-trip time related to the TCP loop, including the propagation delay and the queuing/processing delays suffered by TCP segments in the forward direction and ACK packets in the return direction.

In practical implementations, the upper bound on the TCP window size is set by W_{max}, which is dictated by the availability of socket buffer sizes at the two TCP end-points and application-specific configuration. The throughput equation is then written as

$$thr = \min\left(\frac{W_{\max} \cdot MSS}{RTT}, \frac{k}{\sqrt{p}} \cdot \frac{MSS}{RTT}\right). \tag{5}$$

Eq. 5 suggests that the TCP throughput is basically dependent on (a) packet loss probability and (b) queuing delays since the propagation and processing delay contributions to *RTT* may assumed constant for a given connection path. Since the physical layer in broadband satellite systems is dimensioned to yield bit error rates in the order of 10^{-10}, we may assume $p \rightarrow 0$ for the satellite component of the end-to-end connection. Therefore, the throughput achieved by the satellite components of the TCP connections can be written

$$thr = \frac{W_{\max} \cdot MSS}{RTD + dq_F + dq_R} \tag{6}$$

where dq_F, dq_R are the queuing delays experienced in the forward and return (ACK) path of the TCP connections in the satellite network. *RTD* here refers to the round-trip delay excluding queuing delays.

Eq. 6 suggests that one class of TCP connections may obtain better performance than another in either or under combination of the following cases: a) when they experience lower delays at the MAC scheduler during the resource (i.e. slot) allocation process, b) when the upper bound of the TCP window of their satellite component is set to a higher value. The difficulty arises when we want to control, i.e. quantify, the relative performance differentiation the classes of connections obtain.

In principle, there are several mechanisms and configuration options that can yield the desirable result at the MAC and transport (PEP) layer. We have chosen to deploy the Satellite Waiting Time Priority (SWTP) scheduler for BoD scheduling [8] and vary the W_{max} assigned to TCP connections of different service classes at TCP-PEP.

The SWTP scheduler [8] is an adaptation of the WTP scheduler for satellite networks. SWTP schedules *SR*s from BoD entities rather than individual packets. We have shown that SWTP can provide proportional queuing delay to several classes of MAC frames in the context of the BoD environment. We briefly describe its main aspects below.

Resource request: Formally, if Q_i is the set of newly arrived packets at the queue of BoD entity *i*, i.e. packets that came within the last resource allocation period, *q* the set cardinality, and τ_j the arrival time of packet *j*, $1 \le j \le q$, indexed in increasing order of arrival times, then the BoD entity *m* computes at time *t* the SR timestamp ts_i^m according to the arrival time of the last packet that arrived in the queue during the last resource allocation period, namely: $ts_i = t - \tau_q$.

Resource allocation: The BoD scheduler computes the priority of each SR, $P_i^m(k)$ at k^{th} resource allocation period as

$$P_i^m(k) = r_i^D \cdot \left(w_i^{SR}(k) + \alpha\right) \tag{7}$$

where α accounts for the propagation delay of BTP and the processing delay in BoD entities, while $w_i^{SR}(k) = t - ts_i^m$ and ts_i^m is the timestamp encoded in each *SR*. r_i^D denotes the delay differentiation parameter (DDP). Each one of the *M* MAC layer classes features a specific r_i^D, $1 \le i \le M$.

At the allocation instance, the SWTP allocates TSs by considering requests in decreasing priority order. Requests are fully satisfied as long as they do not exceed the available capacity. All unsatisfied requests will be buffered for the next allocation period. At the next allocation instance, the priorities of the buffered *SR*s will be recalculated to account for the additional waiting time of the request at the scheduler.

At the transport layer, we use W_{max} as the differentiation parameter for the satellite component of the TCP connections. Note that this is different than what several TCP variants, proposed with wireless networks in mind, have done. TCP Peach [9], Westwood [10], and Hybla [11], to mention but a few, actually intervene in the additive increase, multiplicative decrease congestion control mechanism of the protocol. What we do instead is to control W_{max} and this feature exactly makes the scheme independent of the actual TCP variant. W_{max} categorizes persistent TCP connections into two types [3]:

- *Capacity-limited* connection where $W_{max} \ge$ Path Bandwidth-Delay Product, BDP (BDP = RTT.C where C is the bottleneck link capacity).
- *Receive window-limited* connection where $W_{max} <$ Path BDP.

Change of the W_{max} parameter has an impact on TCP throughput in the second case, whereas in the first case the TCP throughput may only be increased with an increase of the bottleneck link capacity.

The split-TCP connections are grouped into *L* transport layer classes, each mapped to a single W_{max} value. Then we can write

$$\frac{W_{\max i}}{W_{\max j}} = \frac{r_i^T}{r_j^T} \quad \forall i, j \in \{1 \dots L\} \tag{8}$$

where $\{r_i^T\}$ is the set of throughput differentiation parameters (TDP). Normalizing with reference to the TDP of the lowest-priority class so that

$$r_1^T > r_2^T > \dots > r_{N-1}^T > r_N^T = 1$$

TCP-PEPs at the border of the network will set the W_{max} of the satellite component of each class i TCP connections to

$$W_{\max(i)} = r_i^T \cdot \min rwnd \tag{9}$$

where min*rwnd* is the minimum value of W_{max} over all classes. For simplicity, we assume that the same number of service differentiation levels is defined at both MAC and transport layers, which equal the total number of TCP traffic classes that are supported within the satellite network; i.e., L equals M and both equal N. Under these assumptions, the problem we face can be stated as follows:

"In a split-TCP capable BoD satellite network, how should one jointly set the TDPs at the TCP-PEPs and the DDPs at the return link MAC scheduler, so that for a given set of DPs, $\{r_i\}$, the PDS model objective of Eq. 3 can be achieved?"

4 Performance Evaluation

4.1 Simulation Setup

We extend the *ns2* to support BoD satellite networks and add an implementation of the SWTP scheduler. The network topology used is shown in Fig. 1, where the bottleneck is assumed to be at the satellite part of the topology. For all simulations, the terrestrial links are configured to be 2048 kbps while the satellite up/downlinks are set to 512 kbps. The packet size is 576 bytes and is fragmented into 48 byte MAC frames. Each TDMA frame period is 24 ms. We use out of band signaling with SRs submitted in pre-assigned slots. The rate granularity is 16 kbps. The effect of our solution is demonstrated by considering persistent TCP sources. Unless explicitly stated otherwise, the network serves three classes with DP set $\{r_i\} \equiv \{1, ½, ¼\}$, i.e., the target throughput ratio between two successively ordered TCP traffic classes is 0.5. If the achieved throughput ratio is smaller than 0.5, then the actual performance spacing is greater than desired and vice versa. We fix min*rwnd* to 8 kb in Eq. 9.

4.2 Impact of Traffic Load on the Two Service Differentiation Mechanisms

We first gain some insight into the impact of each of the two mechanisms separately on TCP throughput.

Transport Layer Differentiation only: We run simulations by switching off the SWTP and then taking a sample TCP connection from each service class for evaluation. Fig. 3 (a) and (b) shows the achieved throughputs under low and high traffic load respectively. The figures suggest that controlling W_{max} at the TCP proxies can yield proportional differentiation only when the system load is low. At low load most of the TCP connections in the network are *receive window-limited* connections and our transport layer configuration suffices to control the differentiation. However, when the network is highly loaded, the TCP connections are no longer limited by their respective W_{max}. Instead, they are *capacity-limited*. Slow start events take place, rendering the receive window constraints inactive. In PDS terminology, the *controllability* and *predictability* of this transport-level differentiation mechanism are highly sensitive to the traffic load. Under low link utilization, the relative differentiation is achieved at steady-state after Slow-Start phase is over.

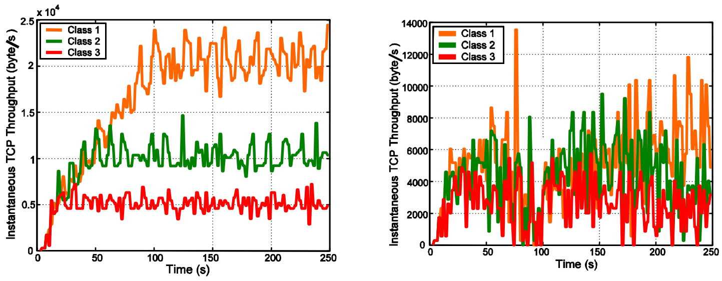

(a) low load *(b) high load*

Fig. 3. Impact of send window control at TCP-PEPs, First-Come First-Served request scheduling at MAC layer

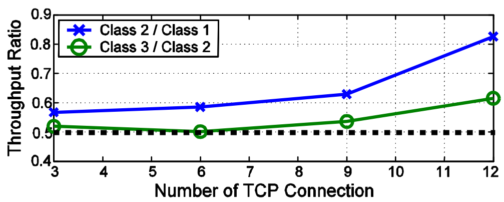

Fig. 4. As the number of TCP connections rises, the PEP loses control of the differentiation

An alternative demonstration of this behavior is shown in Fig. 4 which plots the mean relative performance differentiation versus input traffic load (i.e., number of TCP connections sharing the satellite link). It is clear that as the load increases, the achieved throughput ratios deviate from the target value (0.5). We also see that when the link is saturated, the system practically does not differentiate amongst the three classes, which is unacceptable for satellite operators designing satellite network radio bearer services.

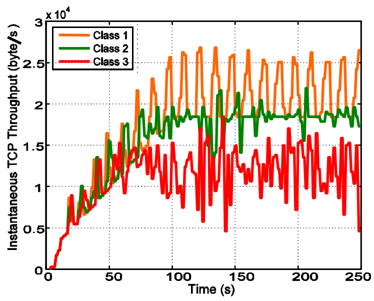

Fig. 5. TCP throughputs differentiated by SWTP; the performance spacing is not proportional

MAC Layer Differentiation only: We then run the simulations with TCP-PEPs switched off while SWTP is switched on. Fig. 5 shows that SWTP alone cannot provide the desired performance spacing even when the load is low. However, the SWTP has been shown in [8] to be increasingly effective when the load is increasing. This property fits in perfectly to complement our transport layer differentiation mechanism, which has been shown to work effectively at low load.

4.3 Evaluation of Full Differentiation Scheme

In Section 4.2, we showed that the two differentiation mechanisms are complementary. By deploying both transport and MAC layer differentiation mechanisms, we have a full service differentiation scheme that works throughout the load range. At low system load, the MAC layer differentiation mechanism is inactive, letting the transport layer differentiation mechanism alone provide the desired performance ratio as needed by the PDS model. As the system load increases, the transport layer differentiation mechanism slowly loses the differentiation control; at the same time, the SWTP scheduler starts taking effect by providing the required additional differentiation, so that the target performance ratio can be maintained.

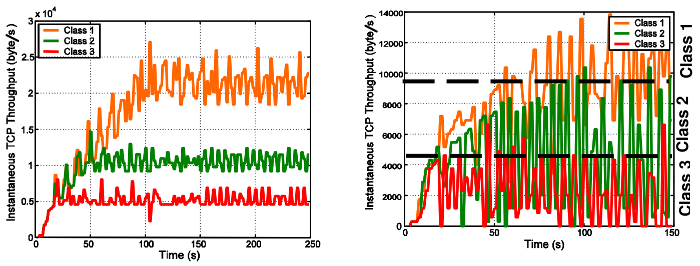

(a) low load *(b) high load*

Fig. 6. Performance of our integrated approach for achieving the PDS model

We first test the integrated scheme with 3 classes of TCP traffic and TDP set = DDP set = {1, ½, ¼}. Fig. 6(a) implies that at low system utilization, the SWTP scheduler does not affect the differentiation provided by the proxy configuration. Fig. 6(b) shows the instantaneous TCP throughput for high load with our integrated approach. It shows improvement when compared with Fig. 3(b), in that now there are distinct throughput levels for each service class. However, there is still considerable fluctuation at individual TCP flow level. We found that the achieved throughput ratios under satellite link saturation are 0.4094 for $\frac{class2}{class1}$ and 0.2793 for $\frac{class3}{class2}$, yielding a total deviation of 0.3113 from the target ratio (i.e. 0.5). This implies that the differentiation given by the combination of SWTP and the proxies has exceeded what is required. We run extensive simulations by varying both DDP and TDP sets. Table 1 shows selected results for the achieved ratios in the form of $\{ \frac{class2}{class1}, \frac{class3}{class2} \}$ throughput ratios. In this case, the optimal settings are DDP set = {1, 0.6, 0.36} and TDP set = {1, 0.6, 0.36} yielding a total deviation of only 0.0595.

Table 1. Achieved TCP throughput ratios between successively ordered classes under different combinations of DDP and TDP sets

DDP Set	TDP Set				
	{1, 0.4, 0.16}	{1, 0.4, 0.2}	{1, 0.5, 0.25}	{1, 0.6, 0.3}	{1, 0.6, 0.36}
{1, 0.4, 0.16}	0.3294, 0.2605	0.3204, 0.3223	0.3648, 0.2230	0.3883, 0.3131	0.3888, 0.2991
{1, 0.4, 0.2}	0.3168, 0.2548	0.2771, 0.3991	0.3574, 0.3008	0.3787, 0.2314	0.3772, 0.3202
{1, 0.4, 0.24}	0.2735, 0.3743	0.2810, 0.4354	0.3536, 0.4148	0.3779, 0.4135	0.3823, 0.3502
{1, 0.4, 0.32}	0.3213, 0.3834	0.3159, 0.4505	0.3510, 0.5050	0.3829, 0.5186	0.3883, 0.5499
{1, 0.5, 0.2}	0.3530, 0.2862	0.3573, 0.3082	0.4149, 0.2372	0.4370, 0.2736	0.4675, 0.2850
{1, 0.5, 0.25}	0.3123, 0.3144	0.3606, 0.3215	0.4094, 0.2793	0.4249, 0.2446	0.4366, 0.3455
{1, 0.5, 0.3}	0.3446, 0.3700	0.3139, 0.4633	0.3276, 0.5491	0.4358, 0.4298	0.3287, 0.5992
{1, 0.5, 0.4}	0.3670, 0.3443	0.3726, 0.4267	0.4060, 0.5062	0.2831, 0.7912	0.7132, 0.5489
{1, 0.6, 0.24}	0.4077, 0.2910	0.3996, 0.3606	0.4823, 0.2816	0.4909, 0.4067	0.3976, 0.4616
{1, 0.6, 0.3}	0.4283, 0.3042	0.4271, 0.3636	0.4765, 0.3888	0.4097, 0.4387	0.5378, 0.3396
{1, 0.6, 0.36}	0.4116, 0.3312	0.4137, 0.3667	0.3728, 0.5668	0.5401, 0.4422	0.5365, 0.5230
{1, 0.6, 0.48}	0.4472, 0.3333	0.4199, 0.3947	0.4845, 0.4534	0.5561, 0.3978	0.5620, 0.4897

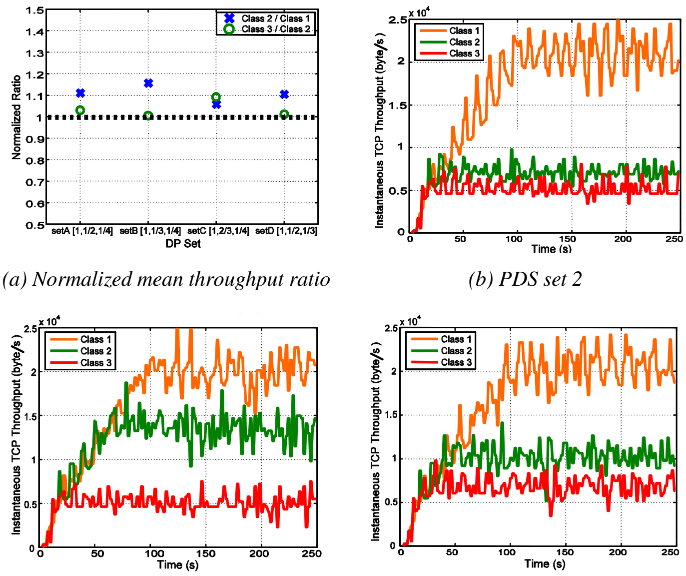

(a) Normalized mean throughput ratio *(b) PDS set 2*

(c) PDS set 3 *(d) PDS set 4*

Fig. 7. Normalized mean and instantaneous throughput ratios under different DP sets

To assess further the capability of our integrated approach to accurately control the spacing between service classes, four sets of DPs have been defined: Set A = {1, ½, ¼}, Set B = {1, 1/3, ¼}, Set C = {1, 2/3, ¼}, Set D = {1, ½, 1/3}. Four sets of simulation runs have been conducted based on the DP sets above with DP = TDP = DDP. Fig. 7(a) plots the *normalized* throughput ratios, i.e., the actual throughput ratios divided over the respective target ratios, for all four cases. The ideal value for the normalized throughput ratio is unity. Whereas, Figures 6(a) and 7(b–d) depict the instantaneous throughput of TCP flows of different service classes under the four DP sets. It can be concluded that our approach allows much flexibility in the control of the spacing between the service classes.

5 Related Work

The provision of relative service differentiation to TCP flows has been primarily investigated in the context of wired networks. In [12], the authors rely on traffic metering, queue management and scheduling mechanisms to do so. The differentiation objective is achieved via marking algorithms used in tandem with explicit congestion notification (ECN) for regulating TCP traffic in the context of class-based service differentiation in [13] and via exploiting the receiver's advertised window (*rwnd*) of TCP connections in [14], where a weighted proportional window control mechanism is proposed. However, since the proposals above are mainly designed for wired network, they either implicitly or explicitly assume that the link capacities are constant. This is not the case for satellite and wireless networks.

Regarding support of QoS over satellite networks, there have been studies mainly on the implementation of the *DiffServ* framework. In [15], the authors assume a fully-fledged ATM switch onboard with buffer management capacity but do not consider the impact of the satellite MAC layer. A gateway architecture to achieve *DiffServ* for satellite networks via a joint resource management and marking approach is proposed in [16]. Their objectives are to minimize bandwidth wastage while satisfying QoS requirements. In [17], the authors compare several buffer management policies for satellite onboard switching to differentiate real time and non-real time traffic.

To the best of our knowledge, this is the first study on addressing *relative* service differentiation for TCP flows in BoD satellite environments. Previously in [18], we look at how to provide relative throughput differentiation for TCP flows in both GEO and non-GEO satellite networks via the use of congestion pricing. Meanwhile, the possibility of the joint use of transport-level and MAC-level service differentiation mechanisms has been demonstrated via analysis and extensive simulations in [19], in the context of split-TCP BoD satellite networks but with Strict Priority MAC level scheduling in mind. The work takes into account the impact of MAC layer in satellite networks when providing service differentiation. We extend this work here by demonstrating that the joint configuration of the two layers can realize a more demanding QoS model over the satellite network, which requires quantitative rather than qualitative relative service differentiation at class level. We not only want to ensure that the performance of one class is better than a lower priority one, we also try to control the performance gap between the two classes.

6 Conclusion

Our paper describes an integrated approach to the provision of proportional throughput differentiation to persistent TCP flows over BoD GEO satellite networks. The approach combines split-TCP proxy agents at transport layer together with a BoD scheduling algorithm, SWTP, exploiting their complementary behavior over the network traffic load range. At low load, the differentiation can be controlled solely by the transport layer differentiation mechanism. At high load, the SWTP scheduling takes over to provide the desired quality spacing between classes, whereas the transport layer mechanism is less effective. With proper configuration of the two mechanisms, one can control the spacing between classes.

The paper contributions are largely methodological. Firstly, it demonstrates that it is feasible to provide differential treatment to TCP flows via combination of transport layer mechanisms and/or MAC layer scheduling algorithms in a BoD environment. Secondly, it describes how to achieve a specific QoS framework (here the PDS model) through our integrated approach.

In the real-world, the transport and MAC layers should be configurable in an automated manner according to the input load. The problem would benefit from analytical methods that can yield the correct parameterization of the two layers for a given traffic mix. We are currently [20] investigating analytical approximations that could assist with this task and we intend to report on our findings in the future.

Acknowledgments. This work is performed within the framework of the SatNEx project, funded by European Commissions (EC) under the Framework Program 6. The financial contribution of the EC towards this project is greatly appreciated.

References

1. C. Dovrolis, D. Stiliadis and P. Ramanathan, “Proportional differentiated services: delay differentiation and packet scheduling,” *IEEE/ACM Transactions on Networking*, vol. 10, no. 1, pp. 12-26, Feb. 2002
2. Y. Zhang, “A multilayer IP security protocol for TCP performance enhancement in wireless networks,” *IEEE JSAC*, vol. 22, no. 4, pp. 767-776, May 2004
3. M. Karaliopoulos, “Support of elastic TCP traffic over broadband geostationary satellite networks,” PhD. Dissertation, University of Surrey, Feb. 2004
4. EN 301 790 V1.3.1, “Digital Video Broadcasting (DVB); Interaction channel for satellite distribution systems,” ETSI European Standard
5. G. Açar, “End-to-end resource management in geostationary satellite networks,” PhD. Dissertation, University of London, Nov 2001
6. M. Mathis et al., “The microscopic behaviour of the TCP congestion avoidance algorithm,” *Computer Communications*, vol. 3, July 1999
7. J. Padhye et al. “Modelling TCP throughput: a simple model and its empirical validation,” in *Proc. ACM SIGCOMM,* Vancouver, 1998
8. W. K. Chai, M. Karaliopoulos and G. Pavlou, “Scheduling for proportional differentiated service provision in geostationary bandwidth on demand satellite networks,” *Proc. of IEEE GLOBECOM,* USA, 2005
9. I. Akyildiz, G. Morabito and S. Palazzo, “TCP-Peach: a new congestion avoidance control scheme for satellite IP networks,” *IEEE/ACM Transactions on Networking*, vol. 9, pp. 307-321, 2001
10. R. Wang et al., “TCP with sender-side intelligence to handle dynamic, large, leaky pipes,” *IEEE JSAC*, vol. 23, no. 2, pp. 235-248, Feb 2005
11. C. Caini and R. Firrincieli, “TCP Hybla: a TCP enhancement for heterogeneous networks,” *Int’l Journal of Satellite Communications and Networking*, vol. 22, pp. 547-566, 2004

12. T. Soetens, S. D. Cnodder and O. Elloumi, “A relative bandwidth differentiated service for TCP micro-flows,” *Proc. of IEEE/ACM Int’l Symposium on Cluster Computing and the Grid*, pp. 602-609, 2001
13. N. Christin and J. Liebeherr, “Marking Algorithms for Service Differentiation of TCP Traffic,” in *Computer Communications, Special Issue on End-to-End Quality of Service Differentiation*, 2005
14. J. Aweya et al., “Weighted proportional window control of TCP traffic,” *Int’l Journal of Network Management*, vol. 11, pp. 213-242, 2001
15. A. Durresi et al., "Achieving QoS for TCP Traffic in Satellite Networks with Differentiated Services," Space Communications, Vol. 17, 2001
16. 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, pp. 351-366, 2003
17. N. Courville, “QoS-oriented traffic management in multimedia satellite systems,” *Int’l Journal of Satellite Communications and Networking*, no. 21, 2003, pp. 367-399
18. W. K. Chai, K.-H. Ho and G. Pavlou, "Achieving Relative Differentiated Services for TCP Traffic over Satellite IP Networks through Congestion Pricing", *Proc. of* NEW2AN, Russia, 2006
19. M. Karaliopoulos et al., “Providing differentiated service to TCP flows over bandwidth on demand geostationary satellite networks,” *IEEE JSAC,* vol. 22, no.2, pp. 333-347, Feb 2004
20. W. K. Chai, M. Karaliopoulos and G. Pavlou, “Proportional Differentiated Service for TCP Flows over Geostationary Bandwidth on Demand Satellite Networks”, Technical report, University of Surrey. Available at http://www.ee.surrey.ac.uk/Personal/W.Chai/publications