

# TCP Connections via Satellite: Cross-Layer Bandwidth Allocation, Pricing and Adaptive Control<sup>‡</sup>

Nedo Celandroni<sup>°</sup>, Franco Davoli\*, Erina Ferro<sup>°</sup>, Alberto Gotta<sup>°</sup>

<sup>°</sup>ISTI-CNR, Area della Ricerca del C.N.R., Via Moruzzi 1, I-56124 Pisa, Italy

[nedo.celandroni@isti.cnr.it](mailto:nedo.celandroni@isti.cnr.it), [erina.ferro@isti.cnr.it](mailto:erina.ferro@isti.cnr.it), [alberto.gotta@isti.cnr.it](mailto:alberto.gotta@isti.cnr.it).

\*Italian National Consortium for Telecommunications (CNIT)

University of Genoa Research Unit, Via Opera Pia 13, 16145 Genova, Italy

National Laboratory for Multimedia Communications, Via Diocleziano 328, Napoli, Italy

[franco.davoli@cnit.it](mailto:franco.davoli@cnit.it).

*Abstract.* The paper focuses on the assignment of a common bandwidth resource to TCP connections over a satellite channel. The connections are grouped according to their source-destination pairs, which correspond to the up- and down-link channels traversed, and each group may experience different fading conditions. By exploiting the tradeoff between bandwidth and channel redundancy (as determined by bit and coding rates) in the maximization of TCP goodput, an overall optimization problem is constructed, which can be solved by numerical techniques. Different relations between goodput maximization and fairness of the allocations are investigated, and a possible pricing scheme is proposed. The allocation strategies are tested and compared in a fading environment, first under static conditions, and then in a real dynamic environment. The goodput-fairness optimization allows significant gains over bandwidth allocations only aimed at keeping the channel Bit Error Rate below a given threshold in all fading conditions.

*Index Terms:* satellite network, resource allocation, optimization, TCP connections.

## I. INTRODUCTION

It is well known to anyone working in satellite communications that the variability in the characteristics of the satellite channel, due to variable traffic loads and to weather conditions that affect the signal attenuation (rain fading), is the main problem to be faced, together with the large propagation delay. Adaptive network management and control algorithms are therefore necessary to maintain the Quality of Service (QoS) of data transmitted over AWGN (Additive White Gaussian Noise) links with high bandwidth-delay product. There is, indeed, a vast literature on performance

---

<sup>‡</sup> Work funded by the Ministry of Education, University and Research (MIUR) in the framework of the “TANGO” project and by the European Commission in the framework of the NoE project “SatNEx” (contract no. 507052).

aspects related to the adaptation of TCP congestion control mechanisms over such channels (see, e.g., [1] and references therein, and [2]), as well as on resource allocation and QoS control in broadband packet networks [3], even in the satellite environment (e.g., [4] and [1, 2]). In particular, cross-layer optimization approaches are becoming widespread for wireless networks in general [5], even though careful design is necessary, to avoid misleading choices [6].

In our study we assume that a number of long-lived TCP connections, also called *elephants* [7], are active on a geostationary satellite network, which consists of  $N$  traffic earth stations, among which a master one, in addition to sending data as any other traffic station, exerts the control on the access to the common resource, i.e., the satellite bandwidth. As the main goal of our study is to investigate the behaviour of TCP connections within a cross-layer resource allocation approach in this type of environment, we do not consider guaranteed bandwidth real-time traffic, whose presence would only introduce dynamic variations (on the time scale of the real-time connections) in the amount of bandwidth available for TCP data. As such, from the point of view of TCP, the presence of real-time traffic can be simply emulated, by a time-varying bandwidth reduction. For the purpose of our study, we then assume that the system has a certain amount of bandwidth to be utilized at its best for TCP traffic. Moreover, for the time being, we do not take into account the presence of short-lived connections (*mice*), whose characterization in this environment presents some additional difficulty, caused by their transient behaviour, and it is currently under investigation. Therefore, our typical configuration consists of fixed numbers of long-lived TCP connections any traffic station (simply “station”) may have with any other in the network. An important assumption that needs to be underlined is that we refer to the TCP Reno congestion control mechanism, and we do not require any modification to the TCP source code. We simply act on the transmission parameters of the satellite link, which are appropriately tuned-up, in order to achieve the maximum end-to-end transfer rate of a TCP connection (also called *goodput*). In [8] and [9], the application of adaptive FEC (Forward Error Correction) techniques was already investigated, to optimize the efficiency of TCP connections when transmitted over rain-faded geostationary satellite channels, with fixed user

antennas, as in the environment that we consider now. We thus adopt the same philosophy and operate at the physical layer as in these previous works, by trading the bandwidth of the satellite link for the packet loss rate due to data corruption. In fact, over wireless links, any gain in the *bit error rate* (BER) (thus, in the packet loss) is generally obtained at the expenses of the *information bit rate* (IBR), and the goodput increases with the IBR and decreases with the BER. The FEC techniques adopted do not interfere in any way with the normal behaviour of the TCP stack, as they are applied just before the transmission over the satellite link. It has been shown in [8] that, given an available radio spectrum, antenna size, and transmission power, the selection of an appropriate modulation scheme and a FEC type allows choosing the BER and the IBR of the link that maximize the goodput of a TCP connection. This optimization can be done for different channel quality conditions, whose variability is due, for example, to the variable attenuation of the signal caused by changing atmospheric events. The optimal transmission parameters, for each channel condition, can be reported in look-up tables and then applied in an adaptive fashion.

Connections take place within different source-destination (SD) pairs over satellite links, which may be generally subject to diverse fading conditions, according to the atmospheric effects at the source and destination stations (i.e., on the up- and down-links). We refer to connections on the same SD pair, which experience a specific channel condition, as belonging to the same “class”; they feed a common buffer at the IP packet level in the traffic station, which “sees” a transmission channel with specific characteristics (that may differ, in general, from those of other SD pairs originating either from the same station or from other stations). The bandwidth allocated to serve such buffers is shared by all TCP connections in that class, and, once fixed, it determines the “best” combination of bit and coding rates for the given channel conditions. The goal of the allocation is to satisfy some global optimality criterion, which may involve *goodput*, *fairness* among the connections, or a combination thereof. Therefore, in correspondence of a specific channel situation, determined by the various up- and down-link fading patterns, and a given traffic load, we face a possible two-criteria optimization problem, whose decision variables are the service rates of the

above mentioned IP buffers for each SD pair, and the corresponding transmission parameters. We will refer to these allocation strategies as TCP-CLARA (Cross Layer Approach for Resource Allocation). Specifically, we consider five allocation criteria within this general philosophy. In all cases, the indexes chosen for the performance evaluation of the system are the TCP connections' *goodput* and the *fairness* of the allocations. The optimal allocations are derived numerically on the basis of an analytical model, under different fade patterns. A possible pricing scheme is also derived, based on the fact that each class of TCP connections "sees" a net bandwidth that is generally obtained at the expense of diverse channel redundancy. The different strategies are compared first in a static fading scenario and then in a dynamically varying one, with fading traces taken from real-life samples. In this case, the allocation is applied adaptively, following the fading and traffic variations.

The paper is organized as follows. The TCP and channel models are defined in the next section. Section III deals with two bandwidth-allocation strategies, based on a mix of goodput and fairness criteria. In Section IV we derive a different optimization criterion, based on Nash Bargaining Solutions [11], and in Section V we give an interpretation in terms of pricing. Section VI defines the physical layer parameters and contains the numerical results and discussion. Conclusions and future developments are stated in Section VII.

## **II. GOODPUT ESTIMATION OF LONG-LIVED TCP CONNECTIONS**

When a number of long-lived TCP sources share the same bottleneck-rate link, it was empirically observed in [12] (by making use of simulation) that, if all connections have the same latency, they obtain an equal share of the link bandwidth. This is strongly supported by our simulations, as well (see Table I below, obtained by using ns-2 [13]), where we suppose the bottleneck is the satellite link. As the latency introduced by a geostationary satellite is quite high (more than half a second), it is reasonable to assume that the additional one introduced by the satellite access network in the

entire path is negligible with respect to the satellite one, and that all connections have the same latency.

In order to avoid time consuming simulations, reasonable estimations can be constructed for the goodput of a TCP Reno agent. A first relation that can be used is the one taken from [14], which is estimated for infinite bottleneck rate, and thus it is valid far apart the approaching of the bottleneck rate itself. Let  $\mu$  be the bottleneck (the satellite link) rate expressed in segments/s,  $n$  the number of TCP sources, and  $\tau$  the delay between the beginning of the transmission of a segment and the reception of the relative acknowledgement, when the satellite link queue is empty. Moreover, let us assume the segment losses to be independent with rate  $q$ . We have  $\tau = c_l + 1/\mu$ , where  $c_l$  is the channel latency. The TCP connections that share the same link also share an IP buffer, inserted ahead of the satellite link, whose capacity is at least equal to the product  $\mu\tau$ . Let also  $b$  be the number of segments acknowledged by each ACK segment received by the sender TCP, and  $T_o$  the timeout estimated by the sender TCP. Then, by exploiting the expression of the send rate derived in [14], dividing it by  $\frac{\mu}{n}$  for normalization and multiplying it by  $1 - q$  for a better approximation, the relative goodput (normalized to the bottleneck rate) can be expressed as

$$T_g = \frac{1 - q}{\frac{\mu}{n} \left( \sqrt{\frac{2bq}{3}} + T_o \min\left[1, 3\sqrt{\frac{3bq}{8}q(1+32q^2)}\right] \right)} \quad (1)$$

Relation (1) is rather accurate for high values of  $q$ , i.e., far apart the saturation of the bottleneck link. For low values of  $q$ , it is found by simulation that, given a fixed value of  $c_l$ , for fixed values

of the parameter  $y = q \frac{\mu\tau}{n}$ , the goodput has a limited variation with respect to individual variations of the parameters  $q$ ,  $\mu$  and  $n$ . Owing to the high number of simulations needed to verify this observation, a fluid simulator (TGEP) has been employed [15], which was validated by means of ns-2, for values of  $y \ll 1$ . Simulation results have been obtained for the goodput estimation, with a

1% confidence interval at 99% level, over a range of values of  $\frac{\bar{y}}{n}$  between 20 and 300, and  $n$  between 1 and 10. For  $0 \leq y \leq 1$ , goodput values corresponding to the same  $y$  never deviate for more than 8% from their mean. We then interpolated such mean values with a 4-th order polynomial approximating function, whose coefficients have been estimated with the least squared errors technique. Assuming a constant  $c_l$ , equal to 0.6 s (it takes into account half a second of a geostationary satellite double hop, plus some processing time), in the absence of the so-called *Delayed ACKs* option ( $b=1$ ), the polynomial interpolating function results to be

$$T_g = a_0 + a_1 y + a_2 y^2 + a_3 y^3 + a_4 y^4 ; \quad y \leq 1, \quad (2)$$

where  $a_0 = 0.995$ ;  $a_1 = 0.11 [s^3]$ ;  $a_2 = -1.88 [s^6]$ ;  $a_3 = 1.98 [s^9]$ ;  $a_4 = -0.63 [s^{12}]$ . For  $y = 1$ ,  $T_g = 0.575$ . For  $y > 1$ , we adopt relation (1), with  $b = 1$ . The adoption of relation (2) for  $y \leq 1$ , instead of relation (1), with 1 as an upper limit, allows reducing the approximation error, which would otherwise reach 25%, within about 8%.

We assume to operate on an AWGN channel, a reasonable approximation for geostationary satellites and fixed earth stations. The segment loss rate  $q$  can be computed as [8]:

$$q = 1 - (1 - p_e / l_s)^{l_s}, \quad (3)$$

where  $p_e$  is the BER,  $l_s$  is the segment length in bits, and  $l_e$  is the average error burst length (*ebL*).

We took  $p_e$  data from the Qualcomm Viterbi decoder data sheet [16] (standard NASA 1/2 rate with constraint length 7 and derived punctured codes), while  $l_e$  was obtained through numerical simulation in [8]. Since we needed to evaluate the BER and error burst characteristics for BER values even less than  $10^{-9}$ , we resorted to extrapolation for some points. The complete set of data is plotted versus  $E_c / N_0$  (channel bit energy to one-sided noise spectral density ratio) in Figure 1 [8].

In order to make  $q$  computations easier, we interpolated such data, and expressed  $p_e$  and  $l_e$  analytically as functions of the coding rate and the  $E_c / N_0$  ratio. We have:

$$\begin{aligned}
p_e(1/2) &= 10^{\square(1.6E_c/N_0+3)} ; & 0 \square E_c/N_0 \square 5 \text{ dB} \\
p_e(3/4) &= 10^{\square(1.6E_c/N_0\square 2.04)} ; & 4 \square E_c/N_0 \square 8 \text{ dB} \\
p_e(7/8) &= 10^{\square(1.6E_c/N_0\square 5)} ; & 6 \square E_c/N_0 \square 10 \text{ dB} \\
l_e(1/2) &= e^{\square 0.32E_c/N_0+1.87} ; & 0 \square E_c/N_0 \square 5 \text{ dB} \\
l_e(3/4) &= e^{\square 0.4E_c/N_0+3.45} ; & 4 \square E_c/N_0 \square 8 \text{ dB} \\
l_e(7/8) &= 63.6\square 19(E_c/N_0)+1.94(E_c/N_0)^2 \square 0.067(E_c/N_0)^3 ; & 6 \square E_c/N_0 \square 10 \text{ dB}
\end{aligned} \tag{4}$$

and, for the uncoded case [17],

$$p_e(1/1) = \frac{1}{2} \operatorname{erfc}\left(10^{(E_c/N_0)/20}\right) ; \quad l_e(1/1) = 1 \tag{5}$$

where  $\operatorname{erfc}(x)$  is the complementary error function  $\frac{2}{\sqrt{\pi}} \int_x^\infty e^{-t^2} dt$ .

The TCP goodput relative to the bottleneck rate is a decreasing function of the segment loss rate  $q$ , which, in its turn, is a decreasing function of the coding redundancy applied in a given channel condition  $C/N_0$  (carrier power to one-sided noise spectral density ratio; see Section VI for the relation between  $C/N_0$  and  $E_c/N_0$ ) and for a given bit rate  $R$ . The combination of channel bit rate and coding rate gives rise to a ‘‘redundancy factor’’  $r \geq 1$ , which represents the ratio between the IBR in clear sky and the IBR in the specific working condition.

The absolute goodput of each TCP connection  $\hat{T}_g$  is obtained by multiplying the relative value by the bottleneck rate, i.e.,

$$\hat{T}_g = T_g \frac{\square}{n} = T_g \frac{1}{n} \cdot \frac{B}{r} \tag{6}$$

where  $B$  is the link rate in segments/s in clear sky conditions.

It is shown in [8] that, for a given hardware being employed (modulation scheme/rate, FEC type/rate), a set of transmission parameters maximize the absolute goodput for each channel condition. Given  $B$  and  $C/N_0$  (which results from a given link budget calculation), and for all possible bit rate values, we must compute  $T_g$  for all allowable coding rates. The actual values of the

goodput are obtained as mentioned above; then, the maximum value is selected. The value of  $T_g$  is taken from (1) or (2), and  $q$  is computed with (3) and (4) or (5). Numerical examples for the link budget corresponding to the Eutelsat satellite Hot Bird 6 are given in Section VI.

### III. THE BANDWIDTH ALLOCATION PROBLEM

We assume that our satellite network operates in single-hop, so that the tasks of the master station are limited to resource assignment and synchronization. Note that, in this respect, we may consider a private network, operating on a portion of the total available satellite capacity, which has been assigned to a specific organization and can be managed by it (as a special case, this situation could also represent a service provider, managing the whole satellite capacity).

It should be clear by now that, if we have a number of links across the satellite network, which represent the physical channels between possible SD pairs, and a total satellite capacity available for the traffic stations, there are different ways of assigning it to the links, according to the goal to be achieved. In any case, since the choice of transmission parameters for a given bandwidth and channel condition influences the TCP goodput, optimized assignments should take into account the particular situation of each link, determined by the current fading, as well as its relative traffic load in terms of ongoing TCP connections.

The problem we address in the following is thus the assignment of bandwidth, bit and coding rates to the IP buffers that serve each specific link, given the rain fading conditions and the traffic load of the satellite network. We make the following assumptions.

1. The end-to-end delay of the TCP connections is the same for each station. This means that the TCP connections are opened in the traffic station itself; they do not come from remote sources, where bottlenecks, possibly present in the terrestrial network, may introduce additional random delays. This corresponds to the case of having the users directly connected to the satellite earth station's router or through a local access network.

2. In each station, there is an IP buffer for each SD pair, and we say that the TCP connections that share it belong to the same class; obviously, they experience the same up-link and destinations' downlink conditions. Irrespective of the station they belong to, let  $N(N \gg 1)$  be the total number of connection classes, corresponding to all possible SD pairs. Some of them may experience the same fading, but they are anyway distinguished, as they share different buffers.

3. In defining and solving the optimal bandwidth-redundancy assignment problems, we consider the system in static conditions. In other words, given a certain number of ongoing connections, distributed among a subset  $F$  of SD pairs, characterized by a certain fading attenuation, we find the optimal assignment as if the situation would last forever. For some of the assignment strategies we consider, this static situation has already been numerically investigated in [10]; a numerical analysis will be also reported in Section VI. Clearly, in a dynamic environment, under variable fading conditions and with starting/ending TCP sessions, a possibility is to perform our calculations at each change of parameters. Some results on this were obtained in [18], and will be further developed in Section VI. It is worth noting that this simple form of adaptation, though involving some cross-layer interaction, does indeed maintain a flavor of separation principle, as it is common in adaptive control and, as such, represents a "cautious" approach, in the sense of [6], not violating any basic layering concept. However, care has to be taken with respect to possible fast fading variations, which might cause oscillations in TCP behavior. This matter will be discussed in the numerical results Section.

We assume that, if the fading conditions of an active class  $i$  ( $i = 1, 2, \dots, F$ ) are such that a minimum goodput  $T_{g,thr}^{(i)}$  cannot be reached by its connections, the specific SD pair would be considered in outage, and no bandwidth would be assigned to it.

Let  $B_i \in [0, W]$  (where  $W$  is the total bandwidth to be allocated, expressed in segments/s),  $r_i \in \{\mathcal{R}^{(1)}, \dots, \mathcal{R}^{(P)}\}$ , and  $n_c^{(i)}$ ,  $i = 1, \dots, F$ , be the bandwidth, the redundancy factor, chosen in the set of available ones (each value  $\mathcal{R}^{(k)}$ ,  $k = 1, \dots, P$ , corresponds to a pair of bit and coding rates), and the

number of connections, respectively, of the  $i$ -th SD pair. Note that, in the cases where different bit and coding rates yield the same redundancy factor, the pair will be selected giving rise to the minimum BER.

The bandwidth assignment consists in the setting of the parameters of each scheduler (which serves the buffers that use a given station's up-link), together with the correspondent optimal channel bit and coding rates. We start by considering two essentially complementary goals that the bandwidth assignment may want to achieve.

- G1) To maximize the global goodput, i.e.,

$$\max_{B_i \in [0, W], r_i \in \{\mathcal{R}^{(1)}, \dots, \mathcal{R}^{(P)}\}_{i=1, \dots, F}} \prod_{j=1}^F n_c^{(j)} \hat{T}_g^{(j)}, \quad (7)$$

$$\text{subject to } \prod_{i=1}^F B_i = W, \quad B_i \geq 0, \quad i = 1, \dots, F \quad (8)$$

- G2) To reach global fairness, i.e., to divide the bandwidth (and to assign the corresponding transmission parameters) in such a way that all TCP connections achieve the same goodput.

Note that, even though the goodput expressions derived in Section II are applied in both cases, the two goals are different and would generally yield different results in the respective parameters: maximizing the global goodput may turn out in an unfair allocation (in the sense that some SD pairs may receive a relatively poor service), whereas, in general, a fair allocation in the above sense does not achieve globally optimal goodput.

As far as the single goals are concerned, the relative calculations may be effected as follows.

- The maximization in (7) is over a sum of separable nonlinear functions (each term in the sum depending only on its specific decision variables, coupled only by the linear equality constraint in (8)). As such, it can be efficiently computed by means of Dynamic Programming [19, 20], if the bandwidth allocations are expressed in discrete steps of a *minimum bandwidth unit (mbu)*, which is the minimum granularity achievable.

- The goodput-equalizing fair allocation can be reached by starting from an allocation proportional to the number of TCP connections per SD pair, computing the average of the corresponding optimal (in the choice of transmission parameters) goodputs, then changing the *mbu* allocations (under constraints (8)) by discrete steps, in the direction that tends to decrease the absolute deviation of each SD pair's goodput from the average, and repeating the operation with the new allocations. A reasonable convergence, within a given tolerance interval, can be obtained in few steps.

As usually one may want to achieve what one believes to be a reasonable combination of goodput and fairness, we propose the following two strategies (termed *Tradeoff* and *Range*, respectively).

**Tradeoff Strategy.** The following steps are performed:

1. Compute the pairs  $(B_i^*, r_i^*)$ ,  $i = 1, \dots, F$ , maximizing the global goodput (7), under constraints (8);
2. Compute the pairs  $(\bar{B}_i, \bar{r}_i)$ ,  $i = 1, \dots, F$ , corresponding to the goodput-equalizing fair choice;
3. Calculate the final allocation as  $\tilde{B}_i = \bar{B}_i \square + B_i^* (1 - \square)$ ,  $i = 1, \dots, F$ , where  $0 < \square < 1$  is a tradeoff parameter, along with the corresponding bit and coding rates.

**Range Strategy.** The following steps are performed:

1. Compute the pairs  $(\bar{B}_i, \bar{r}_i)$ ,  $i = 1, \dots, F$ , corresponding to the goodput-equalizing fair choice;
2. Choose a "range coefficient"  $\square \geq 0$ ;
3. Compute the global goodput-maximizing allocation, by effecting the constrained maximization in (7), with  $\bar{B}_i$  varying in the range  $[\max(\bar{B}_i(1 - \square), 0), \min(\bar{B}_i(1 + \square), W)]$ , instead of  $[0, W]$ ,  $i = 1, \dots, F$ .

As terms of comparison, we will also consider two other possible strategies, termed **BER Threshold** and **Generalized Proportionally Fair (GPF)**, respectively. The first one only assigns the

transmission parameters (bit and coding rates) to each SD pair, in order to keep the BER on the corresponding channel below a given threshold. The bandwidth assignment is then done proportionally to the number of connections per link, multiplied by the corresponding redundancy. The second one is derived from the concept of *Generalized Proportional Fairness* [21], and will be introduced and discussed in Section IV.

In order to comparatively evaluate the different options, we define the following indexes for comparison:

$$\text{Goodput Factor: } \quad \square_g = \frac{\prod_{i=1}^F n_c^{(i)} \hat{T}_g^{(i)}(B_i, r_i)}{\prod_{i=1}^F n_c^{(i)} \hat{T}_g^{(i)}(B_i^*, r_i^*)} \quad (9)$$

where  $(B_i, r_i)$  is a generic choice and  $(B_i^*, r_i^*)$  is the global goodput-maximizing one.

$$\text{Fairness Factor: } \quad \square_f = 1 - \frac{\prod_{j=1}^L |\hat{T}_g^{(j)} - \bar{T}_g|}{2\bar{T}_g(L-1)} \quad (10)$$

where  $L = \prod_{i=1}^F n^{(i)}$  is the total number of ongoing TCP connections, and  $\bar{T}_g = \frac{1}{L} \prod_{k=1}^L \hat{T}_g^{(k)}$  is the average goodput. Note that  $\square_f = 1$  when all goodputs are equal, and  $\square_f = 0$  when the imbalance among the connections' goodputs is maximized, i.e., the goodput is  $\bar{T}_g \cdot L$  for one connection and 0 for the others (yielding a deviation from the average  $(\bar{T}_g \cdot L - \bar{T}_g) + |0 - \bar{T}_g| \cdot (L-1) = 2\bar{T}_g \cdot L - 2\bar{T}_g$ , which is the denominator of (10)).

#### IV. BANDWIDTH-REDUNDANCY OPTIMIZATION IN THE CONTEXT OF NASH BARGAINING PROBLEMS

Different concepts of *fairness* have been used in the literature: *max-min fairness* [22], (as well as the more general concept of *weighted max-min fairness*), *proportional fairness* [23, 24], the *harmonic mean fairness* [25], and the notion of fairness intrinsic in *Nash Bargain Solutions* (NBS)

[26, 11]. All these criteria give rise to Pareto optimal solutions. However, as noted by Kelly [23], max-min fairness gives absolute priority to smaller flows, and does not make efficient use of the available bandwidth. In contrast, proportional fairness allows more resources to be used by non-minimal flows, and provides better efficiency. An assignment is *proportionally fair* if any change in the distribution of the assigned rates would result in the sum of the proportional changes to be non-positive. As we are interested in goodput maximization, we may use a concept of fairness that is defined directly in terms of utilities of the users (the goodputs, in our case), rather than in terms of the bandwidth shares assigned to them, as is done in [21]. This is obtained as the solution of a utility maximization problem, where the performance index to be maximized is the product of the individual utilities, yielding a NBS. Proportional fairness agrees with the NBS in case that the object that is fairly shared is the bandwidth (and the minimum required rate is zero) [21, 24, 11].

Our goal is to find the Nash Bargaining Solution (NBS) with respect to the bandwidth assignment for  $F$  classes of TCP connections, with  $n_c^{(j)}$  connections per class,  $j = 1, \dots, F$ , with initial agreement point  $0$  and performance functions  $f_j^{(i)}(\cdot)$ ,  $i = 1, \dots, n_c^{(j)}$ ,  $j = 1, \dots, F$ , given by the connections' optimal goodputs (in the sense of Section II, i.e., by choosing the redundancy that maximizes the goodput for each given value of bandwidth). More formally:

$$f_j^{(i)}(B_j^{(i)}) = \max_{r_j \in \mathcal{R}^{(1)}, \dots, \mathcal{R}^{(F)}} \hat{T}_g^{(j)}(r_j, B_j^{(i)}) \quad (11)$$

where  $B_j^{(i)}$  is the bandwidth assigned to the  $i$ -th connection in the  $j$ -th class. We assume in the following that the bandwidth assignments can be varied continuously. Functions  $f_j^{(i)}(\cdot)$  are given by the envelope of the goodput curves defined in Section II for each value of redundancy. By defining

$$\bar{B} = [\bar{B}_1, \dots, \bar{B}_F]; \quad \bar{B}_j = [B_j^{(1)}, \dots, B_j^{(n_c^{(j)})}], \quad \text{and} \quad \bar{F} = \sum_{j=1}^F n_c^{(j)}, \quad \text{the set} \quad B_o = \left\{ \bar{B} \in [0, W]^{\bar{F}} : \sum_{j=1}^F \sum_{i=1}^{n_c^{(j)}} B_j^{(i)} = W \right\}$$

is a convex and compact subset of  $\mathbb{R}^{\bar{F}}$ , wherein the individual goodputs satisfy the initial agreement

point performance. If  $f_j^{(i)}(\cdot)$  were concave upper-bounded functions of  $B_j^{(i)}$ , by theorem (2.1) in [11], there would exist a unique NBS, satisfying

$$(P_{\bar{F}}) \quad \max_{\bar{B} \in B_o} \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} f_j^{(i)}(B_j^{(i)}) \quad (12)$$

Moreover, if the functions  $f_j^{(i)}(\cdot)$  are injective on  $B_o$ , by theorem 2.2 in [11], problem  $(P_{\bar{F}})$  is equivalent to

$$(P'_{\bar{F}}) \quad \max_{\bar{B} \in B_o} \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} \ln f_j^{(i)}(B_j^{(i)}) \quad (13)$$

In our situation, the part of functions  $f_j^{(i)}(\cdot)$  that derives from (1) (i.e., for values of  $B_j^{(i)}$  such that the corresponding value of  $y$  (under given parameter values) is greater than 1, can be easily verified to be concave; unfortunately, the same cannot be affirmed in general in the cases of bandwidth values yielding  $0 \leq y \leq 1$ . Therefore, in general, it is not true that the functions are concave in all the range of interest (whereas they are everywhere injective). This does no longer guarantee the uniqueness of the NBS. Nevertheless, it makes sense to investigate the behavior of a cost function like (13), and to analyze the goodput and fairness factors of the corresponding assignments. By discretizing the bandwidth space and applying dynamic programming to solve the problem, as we have done in the previous Section, we would apply anyway a global search method.

The corresponding NBS  $\bar{B}^*$  can be said to be Generalized Proportionally Fair in the sense of [21], i.e.,

$$\prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} \frac{f_j^{(i)}(B_j^{(i)}) \prod f_j^{(i)}(B_j^{(i)*})}{f_j^{(i)}(B_j^{(i)*})} \geq 0 \quad (14)$$

for any other assignment  $\bar{B}$ . We note, in passing, that also the maximization of the sum of the goodputs (rather than the product) would yield a *proportionally fair* solution, in the sense of Kelly [24], or rather, since we are in the presence of a single bottleneck link, a weighted max-min fair one [27, 28], with weights given by the corresponding equilibrium prices (users' willingness to pay).

However, maximizing (13) instead of the sum of the goodputs does have a significant influence on fairness, as implied by the definition of NBS. Indeed, the presence of the logarithm has the effect of changing the equilibrium prices and the corresponding allocations. The interpretation in terms of prices will be given in the next Section.

At this point, we can note that:

- 1) as the TCP connections in each class share the same buffer at the IP packet level, we have no explicit control over their bandwidth utilization, which is determined by the TCP congestion control algorithm, given the rate at which the IP buffer is being served;
- 2) in steady state, the effect of the TCP congestion control algorithm would be that of dividing the bandwidth available for a given class equally among its connections, thus achieving equal goodput for all connections in the same class.

Then, we can do the following, which corresponds, in general, to a sub-optimal assignment. We let the bandwidth of the  $j$ -th class be

$$B_j^* = \sum_{i=1}^{n_c^{(j)}} B_j^{(i)*}, \quad j = 1, \dots, F \quad (15)$$

In the following, we will refer to this as *GPF* or *Proportionally Fair* assignment.

## V. AN INTERPRETATION IN TERMS OF PRICING

Consider the Lagrangian

$$\mathcal{L}(\bar{B}, \lambda) = \sum_{j=1}^F \sum_{i=1}^{n_c^{(j)}} \ln f_j^{(i)}(B_j^{(i)}) + \lambda \left[ W - \sum_{j=1}^F \sum_{i=1}^{n_c^{(j)}} B_j^{(i)} \right] \quad (16)$$

By following the usual interpretation [11], the Lagrange multiplier  $\lambda$  plays the role of *implied cost per unit flow* associated with the (unique) satellite link, whose capacity is being subdivided.

Therefore, we can define the problem associated with the  $i$ -th connection of the  $j$ -th SD pair:

$$\max_{B_j^{(i)} > 0} \left[ \ln f_j^{(i)}(B_j^{(i)}) - \lambda B_j^{(i)} \right], \quad i = 1, \dots, n_c^{(j)}, \quad j = 1, \dots, F \quad (17)$$

and the corresponding total revenue maximization  $\max_{\bar{B} \in B_o} \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} \bar{B}_j^{(i)} B_j^{(i)}$ .

However, the link being a single one,  $\bar{B}_j^{(i)} = \bar{B}$ ,  $\forall i, j$ , and the maximization is effected for

$\prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} B_j^{(i)} = W$ , which simply expresses the requirement of the Pareto optimal solution in this case.<sup>1</sup>

In terms of proportional fairness in the rates per unit charge [24] (which will be referred to in the following as *Weighted Proportional Fairness*), we would have

$$w_j^{(i)*} = \arg \max_{w_j^{(i)} \geq 0} \prod_{j=1}^F \ln f_j^{(i)} \frac{w_j^{(i)}}{\bar{B}} \prod_{j=1}^F w_j^{(i)} \quad (18)$$

for the users' willingness to pay, and, since  $\bar{B}^*$  solves

$$\max_{B_j^{(i)} \geq 0, i=1, \dots, n_c^{(j)}, j=1, \dots, F} \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} w_j^{(i)*} \ln B_j^{(i)}, \quad (19)$$

$$\text{subject to } \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} B_j^{(i)} = W \quad (20)$$

it follows

$$w_j^{(i)*} = \bar{B}_j^{(i)*}; \quad B_j^{(i)*} = \frac{w_j^{(i)*}}{\prod_{k=1}^F \prod_{l=1}^{n_c^{(k)}} w_k^{(l)*}} W \quad (21)$$

Though the conversion to the user utility maximization problem is trivial, it has an interesting interpretation. Indeed, user  $ij$ 's cost of accessing the link can be written as  $\bar{B}_j^{(i)} = \bar{r}_j \bar{q}_j^{(i)}$ , where  $\bar{q}_j^{(i)}$  represents the net bandwidth, and the term  $\bar{r}_j$  can be regarded as the implied cost, *as perceived by all connections of SD-pair  $j$  under the goodput-maximizing redundancy conditions for the given*

<sup>1</sup> We may note, in passing, that the  $\bar{q}_j^{(i)}$ 's we are considering do not correspond to the coefficients of the

weighted sum  $\prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} \bar{q}_j^{(i)} f_j^{(i)}(B_j^{(i)})$ , whose maximization would yield the NBS (which exist, according to Remark 2.4 in [11]). Apparently, the penalty associated with the user problems cannot be a linear function of  $\bar{q}_j^{(i)}$  and  $B_j^{(i)}$ , unless the performance functions are linear in the  $B_j^{(i)}$ 's.

rate. The interpretation then is that a user requiring to operate with a higher redundancy would pay proportionally to it (as this user requires a larger bandwidth share to achieve the same IBR of another). In this respect, it is possible to interpret the redundancies  $r_j^*$ ,  $j = 1, \dots, F$ , as giving rise to  $F$  parallel links, each characterized by a specific cost, and we may regard the net allocations

$$\bar{B}_j = \prod_{i=1}^{n_c^{(j)}} \bar{B}_j^{(i)}, j=1, \dots, F, \text{ as multiple bottlenecks (in the sense that } B_j = r_j^* \prod_{i=1}^{n_c^{(j)}} \bar{B}_j^{(i)} \leq B_j^*).$$

We may then define *Weighted Proportional Fairness* also with respect to the net allocations, as follows. Let  $\bar{B}_j^{(i)*}$  and  $\bar{B}_j^{(i)}$ ,  $i = 1, \dots, n_c^{(j)}$ ,  $j = 1, \dots, F$ , be the optimal allocation and any other allocation, respectively, and  $r_j^*$ ,  $j = 1, \dots, F$ , the goodput-maximizing redundancy corresponding to the optimal allocation. Choose  $\bar{B}_j^{(i)} > 0$  such that  $B_j^{(i)} = r_j^* \bar{B}_j^{(i)}$ ,  $i = 1, \dots, n_c^{(j)}$ ,  $j = 1, \dots, F$ , is feasible (i.e., satisfies constraint (20)). We can then state the following

**Theorem:** the optimal net allocation  $\bar{B}_j^{(i)*}$ ,  $i = 1, \dots, n_c^{(j)}$ ,  $j = 1, \dots, F$ , stemming from the NBS  $\bar{B}^*$ , is Weighted Proportionally Fair (WPF), with respect to link prices  $w_j^{(i)*}$ ,  $i = 1, \dots, n_c^{(j)}$ ,  $j = 1, \dots, F$ .  $\square$

*Proof:* see the Appendix.

If a minimum goodput per connection is imposed, together with an access tariff  $\bar{B}_j$ ,  $j = 1, \dots, F$ , a pricing structure can be defined (analogously to [11]), as

$$p_j^{(i)}(r_j \bar{B}_j^{(i)}) = P_j(\bar{B}_j) + \prod_{i=1}^{n_c^{(j)}} \bar{B}_j^{(i)} \left[ \frac{r_{j,\min}}{r_j} \prod_{i=1}^{n_c^{(j)}} \bar{B}_j^{(i)} \right]^{-1}, \quad i = 1, \dots, n_c^{(j)}, \quad j = 1, \dots, F \quad (22)$$

where  $P(\cdot)$  is a tariff function,  $\bar{B}_{j,\min}^{(i)}$  is the minimum rate per connection necessary to achieve the desired minimum goodput, and  $r_{j,\min}$  is the corresponding goodput-maximizing redundancy. Based on this, a user might be informed of the optimal prices  $r_j^* \bar{B}_j$ , as well as of the cumulative net

bandwidths  $\bar{B}_j^* = \prod_{i=1}^{n_c^{(j)}} \bar{B}_j^{(i)*}$ , related to a given fading condition, and then left free to choose a lowered

price channel, corresponding to a sub-optimal redundancy and net rate. A related admission control and pricing problem could arise here, which is the subject of current research.

## VI. NUMERICAL RESULTS AND DISCUSSION

### A. The satellite network characterization

The fully meshed satellite network considered uses bent-pipe geo-stationary satellite channels and operates in TDMA (Time Division Multiple Access) mode. The master station maintains the system synchronization, other than performing capacity allocation to the traffic stations. The master station performance is the same as the others; thus, any station in the system can assume the role of master. This assures that the master normally operates in pretty good conditions, because when the current master's attenuation exceeds a given threshold, its role is assumed by another station that is in better conditions. To counteract the signal attenuation the system operates bit and coding rates changing. Traffic stations transmit in temporal slots assigned by the master. Table II reports the most significant system parameters. In order to compute the link budget, we considered a portion of the Ka band (20/30 GHz) transponder of the Eutelsat satellite Hot Bird 6, and took data from the file "*Hot Bird 6 data sheet.fm*", which is downloadable from [29]. We consider exploiting 1/4 of the transponder power. Our carrier is modulated in QPSK (quadrature phase shift keying) at 5, 2.5 or 1.25 Msymbols/s; thus, the resulting uncoded bit rates range from 10 to 2.5 Mbit/s. A 1/2 convolutional encoder/Viterbi decoder is employed, together with the punctured 3/4 and 7/8 codes for a possible total 12 combinations of bit/coding rates. The net value of about 7.5 dB of  $E_c/N_0$  ( $C/N_0=77.5$  dBs<sup>-1</sup>), with the maximum modulation rate and the 7/8 coding rate, is assumed as the clear sky condition. In clear sky, after the Viterbi decoder, the bit error rate is about  $10^{-7}$ . The *mbu* size, i.e., the minimum bandwidth unit that can be allocated, has been taken equal to 5 kbit/s; this value is referred to clear sky conditions.

In order to compute the resulting net values of  $E_c/N_0$  at the earth station's receiver input we used the relation

$$E_c / N_0 = C / N_0 \square 10 \text{Log}_{10} R \square m_i, \quad (23)$$

where  $R$  is the uncoded data bit rate in bit/s and  $m_i$  is the modem implementation margin (taken equal to 1 dB). We have assumed  $b=1$  (no *Delayed ACKs* option) and  $T_o = 1.5$  s when using relation (1). We also considered  $l_s = 4608$  bits (576 bytes), which is the default segment length assumed by sender and receiver TCPs, when no other agreement has been possible.

Actually, not all combinations of bit and coding rates must be probed to find the maximum goodput, because some of them result inefficient (i.e., they yield higher BER with the same redundancy). The possible cases are then limited to the following 7 ones: 10 Mbits/s, with code rates 7/8, 3/4, and 1/2; 5 Mbits/s, with code rates 3/4, and 1/2; 2.5 Mbits/s, with code rates 3/4 and 1/2. The uncoded case results inapplicable with the values of  $C/N_0$  available in our situation, even in clear sky conditions.

### ***B. A static case study***

Table III shows the configurations of two of the static tests carried on, denoting the link status ( $C/N_0$  [dB]) and the number of TCP connections in each class. A more complete set of results and comparisons (but regarding only the *Range* and *Tradeoff* strategies) can be found in [10].

Figures 2-3 have been obtained *in static conditions*, namely, with the data in Table III supposed to be holding forever. The software used is the TEAM (*TCP Elephant bandwidth Allocation Method*) free software, which was specifically developed to implement the mechanisms proposed in this paper. The source code is freely available at the web address <http://www.isti.cnr.it/ResearchUnits/Labs/wn-lab/software-tools.html>.

The static results have been validated by means of ns-2 simulations, by running the TCP connections under the bandwidth partitions provided by the TEAM software. The analytical calculations fall within a 99% confidence interval of width less than 1% of the simulation results, and confirm the accuracy of the analytical model adopted for the TCP goodput.

For each record of values in Table III, the figures depict the behavior of the goodput and fairness factors, respectively, for the *Tradeoff*, *Range*, *Generalized Proportionally Fair* (NBS) and *BER*

*Threshold* strategies that have been defined in Section III, for values of the parameters  $\alpha$  and  $\beta$  between 0 and 1. It can be noted that constantly keeping the BER below a given threshold lowers the goodput and does not always maximize the fairness. The Proportionally Fair strategy aims at maximizing the goodput and keeps an optimal value of fairness, but it does not allow changing the fairness value, trading it with the goodput, as is made possible by the *Tradeoff* and *Range* strategies. The expected value of fairness can change due to fading conditions. This effect will be more evident when the results of the dynamic analysis are presented.

The *Tradeoff* and *Range* strategies have a similar behavior, though they span different values of goodput and fairness factors, depending on the system parameters. In all cases, as expected, the goodput factor increases and the fairness factor decreases with increasing  $\alpha$  and  $\beta$ . In general, the span of the *Tradeoff* strategy's goodput and fairness index values is wider in the interval [0, 1] than that of the *Range* strategy, but it must be noted that the parameter  $\beta$  could be increased beyond 1, within the limits imposed by the total bandwidth available. From the obtained results it derives that both strategies allow a sufficient flexibility in choosing a compromise between the two goals of overall goodput maximization and fairness.

### *C. A dynamic case study with real fading*

The simulative analysis reported in the following is the result of the adaptive application of the previously described strategies, when real-life fading attenuation samples (see Figure 4) are dynamically applied. The attenuation data are taken from a real-life data set chosen from the results of the propagation experiment, in Ka band, carried out on the Olympus satellite by the CSTS (Centro Studi sulle Telecomunicazioni Spaziali) Institute, on behalf of the Italian Space Agency (ASI).

The up-link (30 GHz) and down-link (20 GHz) samples considered were 1-second averages, expressed in dB, of the signal power attenuation with respect to clear sky conditions. The attenuation samples were recorded at the Spino d'Adda (North of Italy) station, in September 1992.

We have preferred to use real fading traces, rather than rely upon a model for rain fade generation, as no thoroughly satisfactory model has been devised so far.

The results have been obtained by means of the TEAM software and the ns-2 simulator. The TEAM software calculates both the bandwidth allocations for each class, according to the chosen strategy, and the relative segment loss rates, according to the  $C/N_0$  values of the SD pairs. Ns-2 performs a dynamic simulation, according both to the input trace files, provided by the TEAM software, and the attenuation patterns.

The situation investigated in Figure 5 shows a faded class and a clear-sky class out of ten active ones. Five of the classes are in clear sky condition, while the other five experience different patterns of fading. The total number of TCP connections is 30 and the number of connections per class is reported in Table IV.

Three different allocation strategies are considered: “*Merge*”, *Proportionally Fair (GPF)* and *BER Threshold*, with the threshold set to  $10^{-6}$ . Actually, the *Merge* strategy is not a new one, but simply a merge between the *Range* and the *Tradeoff* ones, in the sense that for each record of the input file the two strategies are separately run for a certain fairness factor threshold (0.85 in the simulation runs); for each record, the allocation values to be passed to ns-2 are then chosen according to the strategy that better performs in terms of goodput factor. Each simulation run gives an observation window of 600 s. In each chart, we trace the behavior of the goodput and of the segment loss rate, respectively, as functions of time.

The allocations of the *Merge* strategy present many oscillations, as is highlighted in Figure 5 (a), in the range 400-600 s. This tendency is due to the optimal choice in a range in which the cost function is quite flat and variations in the bandwidth allocation do not produce a sensible change in the cost value in terms of goodput. However, it must be remarked that the TCP congestion window filters these oscillations, reducing them, as shown in Figure 6.

Figure 7 shows the percentage gain in total goodput, normalized with respect to the worst case, namely, the *BER Threshold* strategy, with the threshold set to  $10^{-7}$ . In the *Range* and *Tradeoff*

strategies, the fairness factor is targeted to 0.85. The absolute goodput value per strategy is given by the sum of the goodputs of all classes, obtained by averaging the values of the dynamic simulation over 600 s. All the allocation strategies are based on an optimization technique, so we do not reasonably expect great differences in performance. Anyway, the merging between *Range* and *Tradeoff*, in the dynamic case, gives a further gain over both the threshold policies, and over the *GPF* one, with a 10% gain with respect to the latter and a gain of 25%-30% with respect to the former ones.

Figure 8 reports the goodput per connection in the three most significant cases considered, for all classes, with the selected fairness factor of 0.85. The *Merge* strategy privileges the faded classes, i.e. classes 6, 7, 8, 9, 10, reducing the allocations and therefore the goodputs of the classes in clear-sky; the allocations and the goodputs are equalized all over the classes also in the presence of hard fading conditions. The goodput gain of the *Merge* strategy for the faded classes is furthermore evident if compared to that of the threshold policy.

As a final comparison, Figure 9 shows that the *Merge* strategy gives a further gain in term of fairness, in comparison with the *GPF* one, over a simulation window. The *Merge* strategy does not experience fading effects, in terms of fairness, as *GPF* instead does, and the fairness factor is kept to the target value.

## VII. CONCLUSIONS

We have considered a problem of cross-layer optimization in bandwidth assignment to TCP connections, traversing different links in a geostationary satellite network, characterized by differentiated levels of fading attenuation. On the basis of the observation that there exists a tradeoff between bandwidth and data redundancy (as determined by bit and coding rate adaptation, used as fade countermeasure at the physical layer) that influences TCP goodput, we have proposed optimization mechanisms that can be used to control the Quality of Service, in terms of goodput and fairness, of the TCP connections sharing the satellite bandwidth. The performance analysis of the

methods proposed, conducted on a few specific cases with real data, by means of the modeling and optimization software developed for this purpose, has shown that relevant gains can be obtained with respect to fade countermeasures that only attempt to constrain the BER below a given threshold, and that a good range of flexibility can be attained in privileging the goals of goodput or fairness. The results have been presented in the dynamic case, where fading levels and number of TCP connections in the system change over time in unpredictable fashion, and can be used in traffic engineering for multiservice satellite networks.

The large amount of simulations done, always with satisfactory results, led us to the conclusion that the adaptive allocation of both bandwidth and redundancy is worth performing, and that all criteria adopted yield relatively comparable results. Current developments are devoted to the further investigation of the relation between some of the allocation methods with pricing schemes, of which a hint has anyway been given in the paper.

## APPENDIX

The condition for  $\varpi_j^{(i)*}$ ,  $i = 1, \dots, n_c^{(j)}$ ,  $j = 1, \dots, F$ , to be WPF, in the sense stated, is the following:

$$\prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} w_j^{(i)*} \frac{\varpi_j^{(i)} \prod \varpi_j^{(i)*}}{\varpi_j^{(i)*}} \geq 0 \quad (A1)$$

Actually, we have

$$\begin{aligned} \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} w_j^{(i)*} \frac{\varpi_j^{(i)} \prod \varpi_j^{(i)*}}{\varpi_j^{(i)*}} &= \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} \varpi_j^* r_j^* \varpi_j^{(i)*} \frac{\varpi_j^{(i)} \prod \varpi_j^{(i)*}}{\varpi_j^{(i)*}} = \\ &= \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} \varpi_j^* r_j^* \left[ \varpi_j^{(i)} \prod \varpi_j^{(i)*} \right] = \varpi_j^* \prod_{j=1}^F r_j^* \varpi_j \prod_{j=1}^F r_j^* \varpi_j^* \prod_{j=1}^F \varpi_j^* = 0 \end{aligned} \quad (A2)$$

As noted by Biddiscombe [30], the equality condition as such is odd, as it would simply mean that  $\varpi_j^{(i)*}$ ,  $i = 1, \dots, n_c^{(j)}$ ,  $j = 1, \dots, F$ , is “no worse” than any other allocation. However, it can be proved in our case that a dual condition to (A1) is satisfied, namely

$$\prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} w_j^{(i)*} \frac{\prod_j^{(i)*} \prod_j^{(i)}}{\prod_j^{(i)}} > 0 \quad (\text{A3})$$

for any other allocation  $\prod_j^{(i)}$ ,  $i=1, \dots, n_c^{(j)}$ ,  $j=1, \dots, F$ , such that  $\prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} B_j^{(i)} = W$ ,

$B_j^{(i)} = r_j^* \prod_j^{(i)}$ ,  $i=1, \dots, n_c^{(j)}$ ,  $j=1, \dots, F$ . Indeed, the right hand side of (A3) can be written as

$$\begin{aligned} \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} r_j^* \prod_j^{(i)*} \frac{\prod_j^{(i)*} \prod_j^{(i)}}{\prod_j^{(i)}} &= \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} r_j^* \prod_j^{(i)*} \frac{\prod_j^{(i)*}}{\prod_j^{(i)}} \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} \prod_j^{(i)} \\ &= \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} B_j^* \frac{\prod_j^{(i)*}}{\prod_j^{(i)}} \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} B_j \frac{B_j^{(i)*}}{B_j^{(i)}} \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} W \end{aligned} \quad (\text{A4})$$

But the function

$$g(\bar{B}) = \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} B_j^* \frac{B_j^{(i)*}}{B_j^{(i)}} \quad (\text{A5})$$

is convex in the domain defined by the constraint set and achieves the minimum value  $W$  in  $\bar{B}^*$ .

This can be easily seen, by writing

$$\frac{\partial g}{\partial B_j^{(i)}} = \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} \frac{B_j^{(i)*}}{B_j^{(i)2}} ; \quad \frac{\partial^2 g}{\partial B_j^{(i)2}} = \frac{2(B_j^{(i)*})^2}{(B_j^{(i)})^3} \quad (\text{A6})$$

and by considering the stationarity conditions of  $g(\bar{B}) + \prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} B_j^{(i)}$ , which are given by

$$\prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} \frac{B_j^{(i)*}}{B_j^{(i)2}} \prod_j^{(i)} = 0, \quad i=1, \dots, n_c^{(j)}, \quad j=1, \dots, F \quad (\text{A7})$$

This implies  $\prod_{j=1}^F \prod_{i=1}^{n_c^{(j)}} B_j^{(i)*} = W \sqrt{\prod_j^{(i)}}$ , i.e.,  $\prod_j^{(i)} = \prod_j^{(i)*}$  and  $B_{j, \min}^{(i)} = B_j^{(i)*}$ ,  $i=1, \dots, n_c^{(j)}$ ,  $j=1, \dots, F$ , so that (A3) is

always satisfied.

## REFERENCES

- [1] S. Kota, M. Marchese, "Quality of Service for satellite IP networks: a survey", *Internat. J. Satell. Commun. Network.*, vol. 21, no. 4-5, pp. 303-348, July-Oct. 2003.

- [2] A. Jamalipour, M. Marchese, H.S. Cruickshank, J. Neale, S.N. Verma (Eds.), Special Issue on “Broadband IP Networks via Satellites – Part I”, *IEEE J. Select. Areas Commun.*, vol. 22, no. 2, Feb. 2004.
- [3] H. J. Chao, X. Guo, *Quality of Service Control in High-Speed Networks*, John Wiley & Sons, New York, NY, 2002.
- [4] F. Alagoz, D. Walters, A. Alrustamani, B. Vojcic, R. Pickholtz, “Adaptive rate control and QoS provisioning in direct broadcast satellite Networks,” *Wireless Networks*, vol. 7, no. 3, pp. 269-281, 2001.
- [5] S. Shakkottai, T. S. Rappaport, P. C. Karlsson, “Cross-layer design for wireless networks”, *IEEE Commun. Mag.*, vol. 41, no. 10, pp. 74-80, Oct. 2003.
- [6] V. Kawadia, P. R. Kumar, “A cautionary perspective on cross layer design”, *IEEE Wireless Commun. Mag.*, 2004 (to appear).  
On-line: [http://black.csl.uiuc.edu/~prkumar/ps\\_files/cross-layer-design.pdf](http://black.csl.uiuc.edu/~prkumar/ps_files/cross-layer-design.pdf).
- [7] M. Ajmone Marsan, M. Garetto, P. Giaccone, E. Leonardi, E. Schiattarella, A. Tarello, “Using partial differential equations to model TCP mice and elephants in large IP networks”, *Proc. IEEE Infocom 04*, Hong Kong, March 2004.
- [8] N. Celandroni, F. Potortì, “Maximising single connection TCP goodput by trading bandwidth for BER”, *Internat. J. Commun. Syst.*, vol. 16, no. 1, pp. 63-79, Feb. 2003.
- [9] N. Celandroni, E. Ferro, F. Potortì, “Goodput optimisation of long-lived TCP connections in a faded satellite channel”, *Proc. 59<sup>th</sup> IEEE Vehic. Technol. Conf.*, Milan, Italy, May 2004.
- [10] N. Celandroni, F. Davoli, E. Ferro, A. Gotta, “A dynamic cross-layer control strategy for resource partitioning in a rain faded satellite channel with long-lived TCP connections”, to be presented at *3<sup>rd</sup> Internat. IFIP Conf. on Network Control and Engineering for QoS, Security and Mobility (Net-Con 2004)*, Palma de Mallorca, Spain, Nov. 2004.
- [11] R. R. Mazumdar, H. Yaïche, C. Rosenberg, “A game theoretic framework for bandwidth allocation and pricing in broadband networks”, *IEEE/ACM Trans. Networking*, vol. 8, no. 5, pp. 667-677, 2000.
- [12] T. V. Lakshman, U. Madhow, “The performance of TCP/IP for networks with high bandwidth-delay products and random loss”, *IEEE/ACM Trans. Networking*, vol. 5, no. 3, pp. 336-350, June 1997.
- [13] The Network Simulator – ns2. Documentation and source code from the home page: <http://www.isi.edu/nsnam/ns>.
- [14] J. Padhye, V. Firoiu, D. F. Towsley, J. F. Kurose, “Modeling TCP Reno performance: a simple model and its empirical validation”, *IEEE/ACM Trans. Networking*, vol. 8, pp. 133-145, 2000.

- [15] <http://wnet.isti.cnr.it/software/tgep.html>.
- [16] “Q1401: K=7 rate 1/2 single-chip Viterbi decoder technical data sheet”, *Qualcomm, Inc.*, Sept. 1987.
- [17] J. G. Proakis, *Digital communications*, McGraw Hill, New York, NY, 1995.
- [18] N. Celandroni, F. Davoli, E. Ferro, A. Gotta, “Adaptive optimal bandwidth partitioning among TCP elephant connections over multiple rain-faded satellite channels”, submitted to *3<sup>rd</sup> Internat. Workshop on QoS in Multiservice IP Networks (QoS-IP 2005)*, Catania, Italy, Feb. 2005.
- [19] K. W. Ross, *Multiservice Loss Models for Broadband Telecommunication Networks*, Springer, London, UK, 1995.
- [20] N. Celandroni, F. Davoli, E. Ferro, “Static and dynamic resource allocation in a multiservice satellite network with fading”, *Internat. J. Satell. Commun. Network.*, vol. 21, no. 4-5, pp. 469-487, July-Oct. 2003.
- [21] C. Touati, E. Altman, J. Galtier, “On fairness in bandwidth allocation”, Research Report no. 4296, INRIA, Sophia Antipolis, France, Sept. 2001.
- [22] D. Bertsekas, R. Gallager, *Data Networks*. Prentice-Hall, Englewood Cliffs, NJ, 1987.
- [23] F. P. Kelly, “Charging and rate control for elastic traffic”, *Europ. Trans. Telecommun.*, vol. 8, pp. 33-37, 1998.
- [24] F. P. Kelly, A. Maulloo, D. Tan, “Rate control in communication networks: shadow prices, proportional fairness and stability”, *J. Oper. Res. Society*, vol. 49, pp. 237-252, 1998.
- [25] L. Massoulié, J. W. Roberts, “Bandwidth sharing and admission control for elastic traffic”, *Telecommun. Syst.*, vol. 15, pp. 185-201, 2000.
- [26] R. Mazumdar, L. G. Mason, C. Douligeris, “Fairness in network optimal flow control: optimality of product forms”, *IEEE Trans. Commun.*, vol. 39, pp. 75-782, May 1991.
- [27] P. Marbach, “Priority service and max-min fairness”, *IEEE/ACM Trans. Networking*, vol. 11, no. 5, pp. 733-746, Oct. 2003.
- [28] P. Marbach, “Analysis of a static pricing scheme for priority services”, *IEEE/ACM Trans. Networking*, vol. 12, no. 2, pp. 312-325, April 2004.
- [29] <http://www.eutelsat.com/satellites/13ehb6.html>
- [30] M. Biddiscombe, “Proportional fairness – a query”, available online from: <http://www.priatel.co.uk/martin/work/propfair.pdf>.

TABLE I. NORMALIZED GOODPUT OF 5 TCP CONNECTIONS WITH CONFIDENCE INTERVALS AT 99% LEVEL FOR DIFFERENT VALUES OF THE SEGMENT LOSS RATE.

TCP Reno (no delayed ACKs) goodput of 5 connections sharing a link with a bottleneck rate of 455 segments/s, link latency 0.5 s - (ns2 simulations)						
$q$	Connect. #1	Connect. #2	Connect. #3	Connect. #4	Connect. #5	Total
$10^{-4}$	0.198±3.1%	0.199±2.5%	0.199±0.5%	0.196±2.2%	0.202±2.3%	0.994±0.03%
$10^{-3}$	0.165±2.4%	0.169±2.6%	0.164±4.3%	0.169±5.3%	0.169±1.3%	0.836±0.67%
$10^{-2}$	0.0469±0.64%	0.0471±2.3%	0.470±1.5%	0.0466±3.1%	0.0468±1.6%	0.234±0.4%
$10^{-1}$	0.0090±3%	0.0090±0.8%	0.0090±0.7%	0.0090±0.6%	0.0092±2%	0.0452±0.4%




Fig. 1. BER and average burst length versus  $E_c / N_0$  for different values of convolutional coding rates. Extrapolated values are circled.

TABLE II. MOST SIGNIFICANT VALUES OF THE TDMA SYSTEM CONSIDERING THE HOT BIRD 6 KA PAYLOAD.

Stations' antenna diameter	1.2 m
Stations' power	7 dBW
Satellite $G/T$	13 dB°K
Satellite transponder $E.I.R.P.$ (effective isotropic radiation power)	52 dBW
Share of the satellite transponder power	1/4
Maximum/minimum capacity of the carrier (QPSK modulation)	10/2.5 [Mbit/s]
Net $E_c / N_0$ in clear sky conditions (10Mbit/s)	7.5 [dB]
Possible data coding rates	7/8 (clear sky), 3/4, 1/2
Information bit rate in clear sky (10Mbit/s at 7/8 coding rate)	8.75 [Mbit/s]
Information bit rate in clear sky after system overhead	8[Mbit/s] = 1600 <i>mbus</i>

TABLE III. CONFIGURATION OF THE 2 TESTS ( $n_c$  is the number of TCP connections in the relevant class).

Record#	Class 1	Class 2	Class 3	Class 4	Class 5	Class 6	Class 7	Class 8	Class 9	Class 10
	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$
1	78.0;2	77.0;3	76.1;2	68.9;4	74.0;2	75.0;3	76.5;1	71.8;4	76.3;3	76.6;6
2	78.0;2	68.0;3	76.1;2	68.9;4	73.0;2	73.0;3	75.0;1	72.0;4	76.3;3	76.6;6

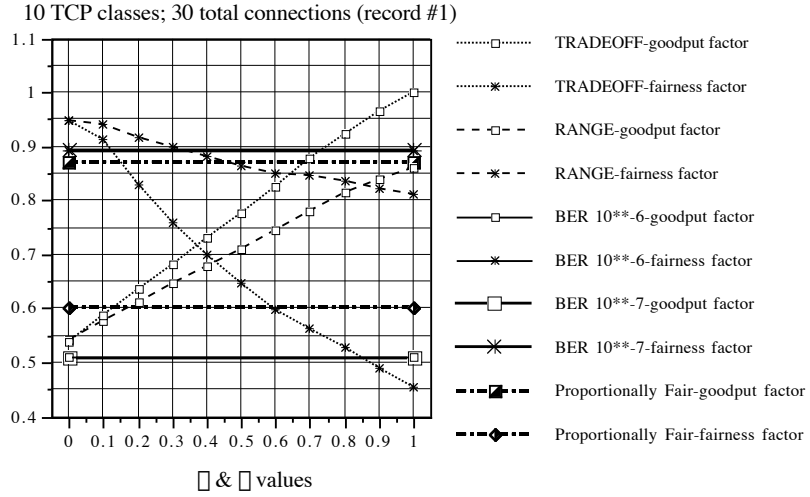


Fig. 2. Goodput and fairness indexes for the data in record #1.

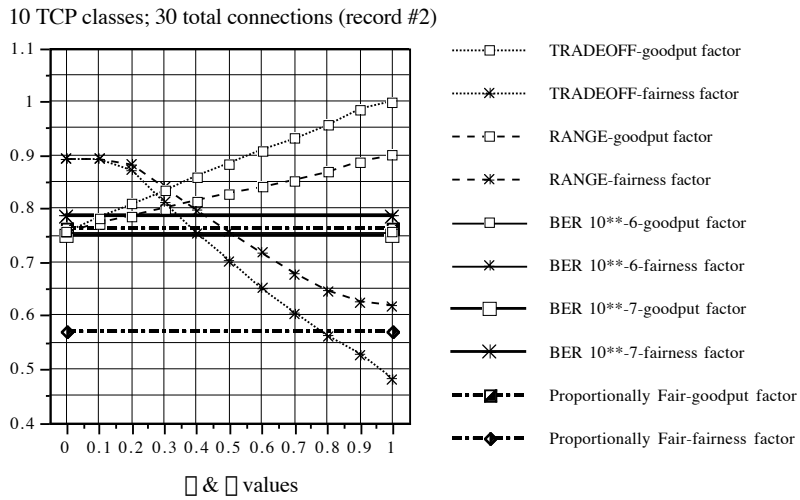


Figure 3. Goodput and fairness indexes for the data in record #2.

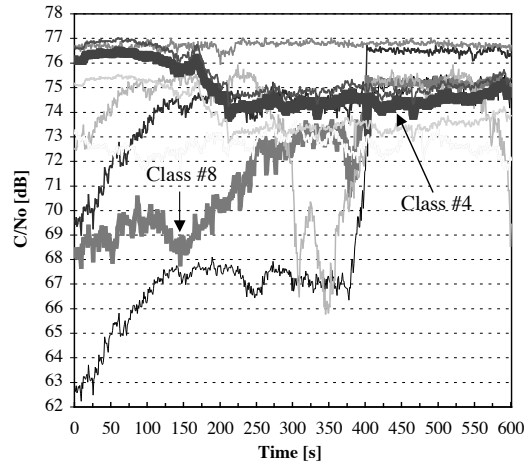




Fig. 4.  $C/N_0$  values vs time represented for the 10 classes.

TABLE IV. NUMBER OF TCP CONNECTIONS PER CLASS


	Class #1	Class #2	Class #3	Class #4	Class #5	Class #6	Class #7	Class #8	Class #9	Class #10
Conn.	2	3	3	2	2	5	3	3	2	4



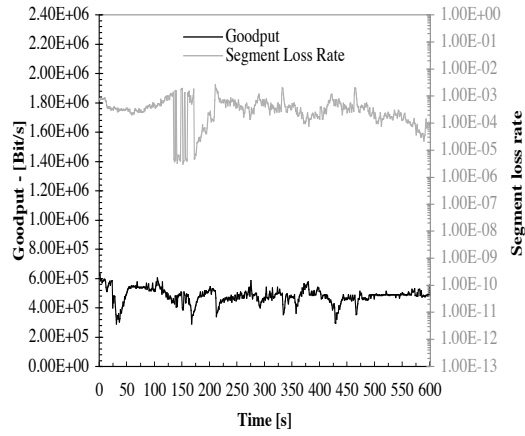
(a) – Merge



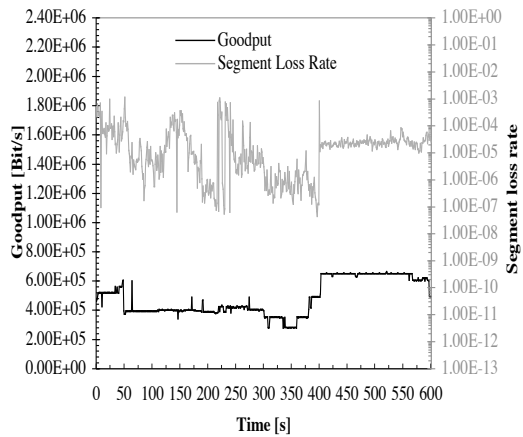
(b) - Merge




(a) – Proportionally Fair



(b) – Proportionally Fair



(a) – BER Threshold



(b) – BER Threshold

Fig 5. Merge, Proportionally Fair and BER Threshold ( $\text{thr}=10^{-6}$ ) strategy for the class #8 in fading (a) and for the class #4 in clear sky (b)




Fig. 6. TCP congestion window of class #8 vs time.

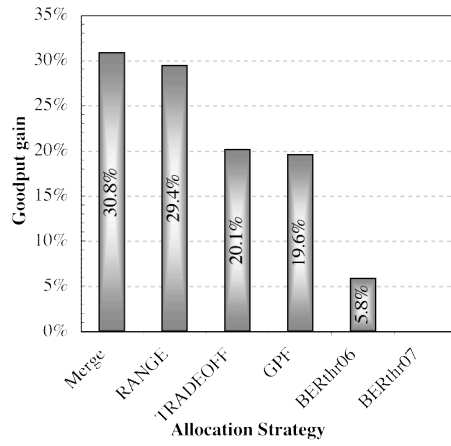


Fig. 7. Total goodput gain averaged over all classes compared with the worst case.

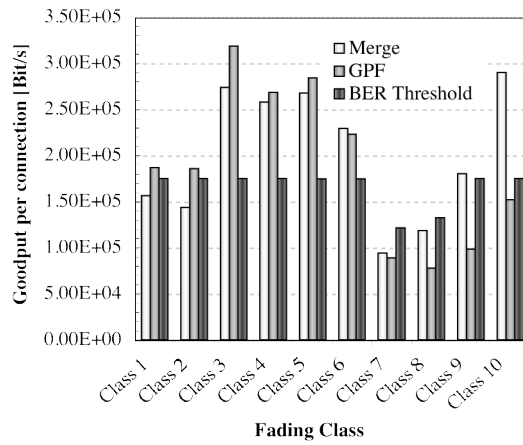


Fig. 8. Goodput per TCP class per connection for the most significant strategies.

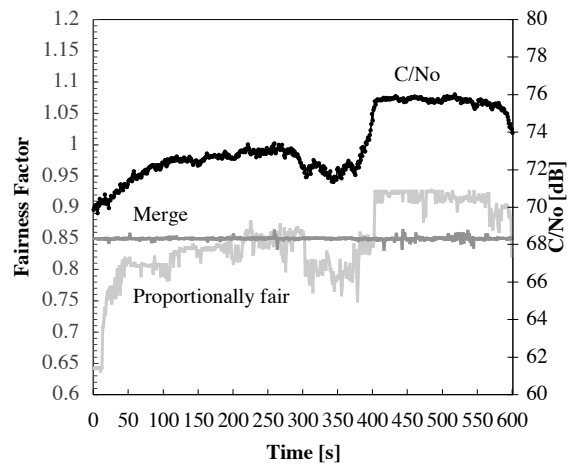


Fig. 9. Fairness factor of *Merge* and *Proportionally Fair* strategy vs time, with an indicative value of  $C/N_0$ , averaged over classes.