

A DYNAMIC CROSS LAYER CONTROL STRATEGY FOR RESOURCE PARTITIONING IN A RAIN FADED SATELLITE CHANNEL WITH LONG-LIVED TCP CONNECTIONS^(*)

Nedo Celandroni¹ Franco Davoli² Erina Ferro¹, Alberto Gotta¹

¹*ISTI-CNR, Area della Ricerca del C.N.R., Via Moruzzi 1, I-56124 Pisa, Italy, {nedo.celandroni,erina.ferro,alberto.gotta}@isti.cnr.it;* ²*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*

Abstract: The paper aims at devising a control system for dynamic resource allocation in a packet-oriented satellite network. The traffic to be served is represented by TCP long-lived connections (elephants). A Master Station adaptively assigns bandwidth and transmission parameters (bit and coding rate) to TCP buffers at the earth stations, grouping connections characterized by the same source-destination pair. The assignment is effected according to each pair's traffic load and fading conditions, in order to reach a common goal. The latter may consist of maximizing the overall TCP goodput, of equalizing the connections' goodput for global fairness, or a combination thereof. Three different allocation strategies are devised, and their respective performance is compared, under a realistic link budget.

Key words: Satellite Communications, Fade Countermeasures, Cross-Layer Optimization, TCP, Dynamic Bandwidth Allocation.

(*) Work supported by MIUR (Ministero dell'Istruzione, Università e Ricerca) in the framework of the "TANGO" and "DIDANET" projects, by the Italian National Research Council (CNR), under the *IS-MANET* project, and by the European Commission in the framework of the *SatNEx* NoE project (contract N. 507052).

1. INTRODUCTION

Satellite systems not only have to face variable load multimedia traffic, but also variable channel conditions with large propagation delay. The variability in operating conditions is due both to changes in the traffic loads and to the signal attenuation on the satellite links, because of bad atmospheric events, which particularly affect transmissions in the Ka band (20-30 GHz). It is therefore stringent to make use of adaptive network management and control algorithms to maintain the Quality of Service (QoS) of the transmitted data. Our study considers the following scenario. A geo-stationary satellite network consists of N stations, among which a master station exerts, in addition, the control on the access to the common resource, i.e., the satellite bandwidth, and $N-1$ traffic stations exchange non-real-time (bulk) traffic. The traffic consists of a number of long-lived TCP connections (also called *elephants*¹) any station may have with any other station in the network. In Celandroni^{2,3} the application of adaptive FEC (Forward Error Correction) was investigated, to optimize the efficiency of TCP connections over AWGN (additive white Gaussian noise) links with high delay-bandwidth product. This case well matches transmissions over rain faded geo-stationary satellite channels, with fixed user antennas. In our study, we thus adopt the same philosophy and operate at the physical level, 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) (i.e., in the packet loss) is generally obtained at the expenses of the *information bit rate* (IBR), and the end-to-end transfer rate of a TCP connection (also called *goodput*) increases with the IBR and decreases with the BER. The adaptation techniques adopted do not interfere in any way with the normal behavior of the TCP stack; the end-to-end protocols are thus left unaltered, while the transmission parameters of the satellite link are appropriately tuned-up.

In this paper, we study bandwidth, bit and coding rate allocation methods for TCP long-lived connections that are in effect on a number of source-destination (SD) pairs over satellite links, which may be generally subject to different fading conditions. In this situation, connections belonging to the same SD pair feed a common buffer at the IP packet level in the earth station, which “sees” a transmission channel with specific characteristics; the latter may generally be different from those of other SD pairs originating from the same station or from other stations. The bandwidth allocated to serve such buffers is shared by all TCP connections in that group, 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

situation of channel conditions, 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).

Though there is a vast literature on performance aspects related to the adaptation of TCP congestion control mechanisms over large bandwidth-delay product and error-prone satellite channels (see, Kota⁴ and references therein and Jamalipour⁵), as well as on resource allocation and Quality of Service (QoS) control in broadband packet networks⁶, even in the satellite environment^{4, 5, 7}, to the best of the authors' knowledge, this is the first attempt to conjugate TCP over satellite performance and resource allocation, in the framework of a cross-layer approach.

2. 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 Lakshman⁸ (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 (obtained by using Network Simulator⁹ – ns2), where we suppose the bottleneck is the satellite link. As the latency introduced by a geo-stationary satellite is quite high (more than half a second) it is reasonable to assume that the additional latency introduced by the satellite access network in the entire link 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 Padhye¹⁰, which is estimated for infinite bottleneck rate, and thus it is valid far apart the approaching of the bottleneck rate itself. Let μ be the bottleneck (the satellite link) rate expressed in segments/s, n the number of TCP sources, and τ 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, 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 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 Padhye¹⁰, dividing it by $\frac{\mu}{n}$ for normalization and multiplying it by $1-q$ for a better approximation, the relative (normalized to the bottleneck rate) goodput can be expressed as

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

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 $y = q \left(\frac{\mu}{n} \right)^2 \tau^5$, the goodput has a limited variation with respect to individual variations of the parameters q , μ and n . Owing to the high number of simulations needed to verify this observation, a fluid simulator¹¹ has been employed, which was validated by means of ns2, for values of $y \leq 1$. Simulation results have been obtained for the goodput estimation, with a 1% confidence interval at 99% level, over a range of values of $\frac{\mu}{n}$ 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 (a value that takes into account half a second of a geostationary satellite double hop, plus some delay for terrestrial links and 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 instead relation (1), with $b = 1$.

We assume to operate on an AWGN channel, a reasonable approximation for geo-stationary satellites and fixed earth stations. The segment loss rate q can be computed as in Celandroni²:

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

where p_e is the bit error probability (BER), l_s is the segment bit length and l_e is the average error burst length (*ebf*). We took p_e data from the Qualcomm Viterbi decoder data sheet¹² (standard NASA 1/2 rate with constraint length 7 and derived punctured codes), while l_e was obtained through numerical simulation in Celandroni². (channel bit energy to one-sided noise spectral density ratio).

In order to make q computations easier, we interpolated data taken from Celandroni² and expressed p_e and l_e analytically as functions of the coding rate and the E_c / N_0 (channel bit energy to one-sided noise spectral density ratio) ratio. We have

$$\begin{aligned} p_e(1/2) &= 10^{(1.6E_c/N_0+3)}; & 0 \leq E_c/N_0 \leq 5 \text{ dB} \\ p_e(3/4) &= 10^{(1.6E_c/N_0-2.04)}; & 4 \leq E_c/N_0 \leq 8 \text{ dB} \\ p_e(7/8) &= 10^{(1.6E_c/N_0-5)}; & 6 \leq E_c/N_0 \leq 10 \text{ dB} \\ l_e(1/2) &= e^{-0.32E_c/N_0+1.87}; & 0 \leq E_c/N_0 \leq 5 \text{ dB} \\ l_e(3/4) &= e^{-0.4E_c/N_0+3.45}; & 4 \leq E_c/N_0 \leq 8 \text{ dB} \\ l_e(7/8) &= 636 - 19(E_c/N_0) + 1.94(E_c/N_0)^2 - 0.067(E_c/N_0)^3; & 6 \leq E_c/N_0 \leq 10 \text{ dB} \end{aligned} \quad (4)$$

For the uncoded case¹³ we have

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

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 4 for the relation between C/N_0 and E_c/N_0) and for a given bit rate b_r . The combination of channel bit rate and coding rate gives rise to a “redundancy factor” $r_{cs} \geq 1$, 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.,

$$\hat{T}_g = T_g \frac{\mu}{n} = T_g \frac{1}{n} \cdot \frac{B}{r_{cs}} \quad (6)$$

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

In Celandroni² it is shown 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 rates, 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 4.

3. THE BANDWIDTH ALLOCATION PROBLEM

Consider now a satellite network, consisting of a bent-pipe (transparent) geo-stationary payload and $N-1$ traffic stations. We assume to operate 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 might also represent a service provider, managing the whole satellite capacity).

The problem we address in the following is the assignment of bandwidth, bit and coding rates to the IP buffers serving each specific link, given the fading conditions and the load of the network.

We make the following assumptions.

1. The end-to-end delay of the TCP connections is the same for each station.
2. In each station, there is an IP buffer for each SD pair, and we say that the TCP connections sharing it belong to the same class; obviously, they experience the same up-link and destinations' downlink conditions.
3. In the present paper, we consider the bandwidth assignment 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. Possible ways of adaptive allocation in a dynamic environment are the matter of ongoing investigation.

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), $r_{cs}^{(i)} \in \{R_1, \dots, R_P\}$, and $n_c^{(i)}$, $i = 1, \dots, F$, be the bandwidth, the redundancy factor, chosen in the set of available ones (each value R_i 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 that gives rise to the minimum BER.

We consider two basically opposite ways of assigning the bandwidth (which corresponds to setting the parameters of the scheduler serving the buffers that use a given station's up-link), together with the transmission parameters:

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

$$\max_{B_i \in [0, W], r_{cs}^{(i)} \in \{R_1, \dots, R_P\}, i=1, \dots, F} \sum_{j=1}^F n_c^{(j)} \bar{P}_g^{(j)}, \quad (7)$$

$$\text{subject to } \sum_{i=1}^F B_i = W \quad (8)$$

G2) To reach global fairness, i.e., to divide the bandwidth (and assign the corresponding transmission parameters, namely, channel bit and coding rate) in such a way that all SD pairs achieve the same goodput.

Note that, even though the goodput optimization formulas of Section 2 are applied in both cases, the two goals are different and will generally yield different results in the respective parameters: maximizing the global goodput may result in an unfair allocation (in the sense that some SD pairs may receive a relatively poor service), whereas a fair allocation generally 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 constraint (8)). As such, it can be computed efficiently by means of Dynamic Programming^{14, 15}, if the bandwidth allocations are expressed in discrete steps of a *minimum bandwidth unit (mbu)*, which is the minimum granularity that can be achieved.

As regards the goodput equalizing fair allocation, it can be reached by starting from an allocation proportional to the number of TCP connections per SD pair, by computing the average of the corresponding optimal (in the

choice of transmission parameters) goodputs, then changing the *mbu* allocations (under constraint (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, which will then be evaluated numerically in the next Section.

Tradeoff Strategy. The following steps are performed:

Compute the pairs (B_i^*, r_i^*) $i=1, \dots, F$, maximizing the global goodput (7), under constraint (8);

Compute the pairs (\bar{B}_i, \bar{r}_i) $i=1, \dots, F$, corresponding to the goodput equalizing fair choice;

Calculate the final allocation as $\tilde{B}_i = \bar{B}_i \rho + B_i^* (1 - \rho)$, $i=1, \dots, F$, where $0 \leq \rho \leq 1$ is a tradeoff parameter, along with the corresponding bit and coding rates.

Range Strategy. The following steps are performed:

Compute the pairs (\bar{B}_i, \bar{r}_i) $i=1, \dots, F$, corresponding to the goodput equalizing fair choice;

Choose a "range coefficient" $\beta \geq 0$;

Compute the global goodput maximizing allocation, by effecting the constrained maximization in (7), with \bar{B}_i varying in the range $[\max(\bar{B}_i(1 - \beta), 0), \min(\bar{B}_i(1 + \beta), W)]$, instead of $[0, W]$, $i=1, \dots, F$.

As a term of comparison, we will also consider another possible strategy (termed *BER Threshold* in the following), which only assigns the transmission parameters (bit and coding rate) to each SD pair, in order to keep the BER on the corresponding channel below a given threshold. The bandwidth assignment is done proportionally to the number of connections per link, multiplied by the corresponding redundancy.

In order to comparatively evaluate the different options, we define the following terms of comparison:

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

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

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

where $L = \sum_{i=1}^F n^{(i)}$ is the total number of ongoing TCP connections, and

$\bar{T}_g = \frac{1}{L} \sum_{k=1}^L \hat{T}_g^{(k)}$ is the average goodput. Note that $f_f = 1$ when all goodputs are equal, and $f_f = 0$ when the unbalance 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)).

4. NUMERICAL RESULTS

We considered a fully meshed satellite network that uses bent-pipe geostationary satellite channels. This means that the satellite only performs the function of a repeater and it does not make any demodulation of data. The system operates in TDMA 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, the role of master can be assumed by any station in the system. 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.

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 the web¹⁶.

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 ($C/N_0=77.5$ dB), 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 relation (11) below.

$$E_c/N_0 = C/N_0 - 10\text{Log}_{10}b_r - m_i, \quad (11)$$

where b_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 $l_s = 4608$ (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.

Table 1. Configuration of the 3 tests (n_c is the number of TCP connections in the relevant class).

Record #	Class 1	Class 2	Class 3	Class 4	Class 5
	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$	$C/N_0;n_c$
1	78.0;2	77.0;3	76.1;2	68.9;4	74.0;2
2	78.0;2	68.0;3	76.1;2	68.9;4	73.0;2
3	78.0;2	72.0;3	76.1;2	69.9;4	73.0;2
Record #	Class 6	Class 7	Class 8	Class 9	Class 10
1	75.0;3	76.5;1	71.8;4	76.3;3	76.6;6
2	73.0;3	75.0;1	72.0;4	76.3;3	76.6;6
3	75.0;3	76.5;1	71.8;4	76.3;3	63.0;6

Table I shows the configurations of the three tests carried on, denoting the link status (C/N_0 [dB]) and the number of TCP connections in each class.

The following Figs 1 and 2 have been obtained, by using the above data, by means of the TEAM (*TCP Elephant bandwidth Allocation Method*) modeling and optimization software, which was specifically developed to implement the mechanisms proposed in this paper and other, more general ones, which will be the subject of forthcoming research.

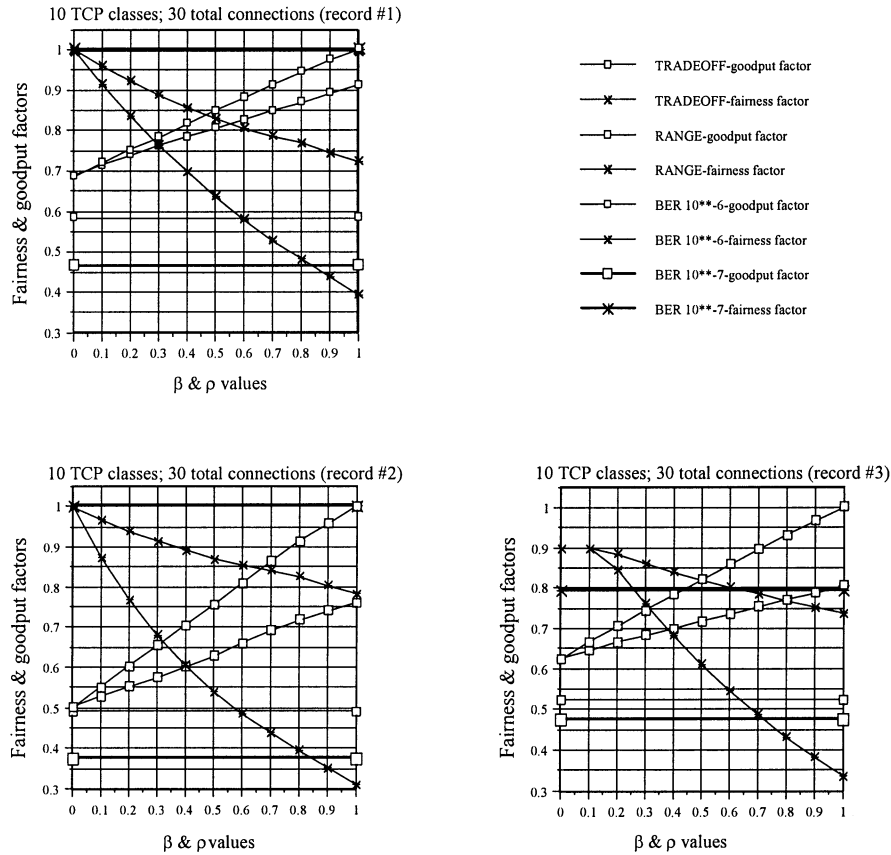


Figure 1 Goodput and fairness indexes for the data in record #1 #2 #3.

For each record of values in Table I, the figures depict the behavior of the goodput and fairness factors, respectively, for the *Tradeoff*, *Range*, and *BER Threshold* strategies that have been defined in the previous Section, for values of the parameters ρ and β between 0 and 1. It can be noted that constantly keeping the BER below a given threshold lowers the goodput and not always maximizes the fairness (as can be seen in Fig. 1 record #3); moreover, enforcing a constant lower BER (10^{-7}) lowers the goodput, without any appreciable gain in fairness.

The *Tradeoff* and *Range* strategies have a similar behavior, though they span different values of goodput and fairness factors, depending on the sys-

tem parameters. In all cases, as expected, the goodput factor increases and the fairness factor decreases with increasing ρ and β . 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 β could be increased beyond 1, within the limits imposed by the total bandwidth available. A conclusion that can be drawn from the obtained results is that both strategies allow a sufficient flexibility in choosing a compromise between the two goals of overall goodput maximization and fairness.

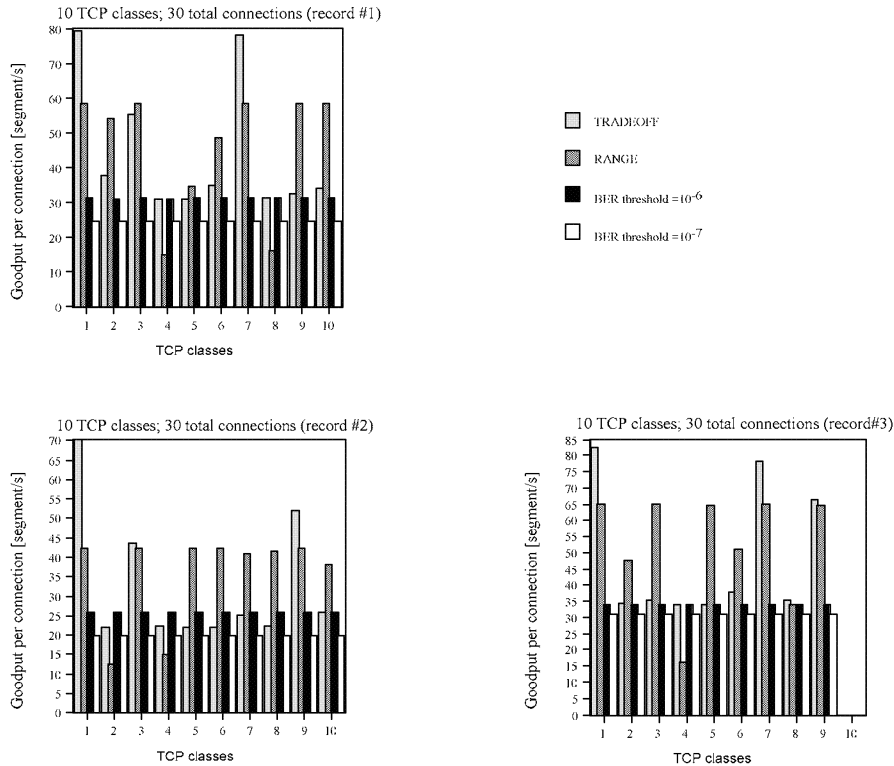


Figure 2. Goodput per TCP class for the data in record #1 #2 #3. (class 10 of record #3 is in outage for all strategies).

The total goodput of the system [segments/s] calculated by the TEAM modeling and optimization software, is in total agreement with simulation results of ns2. The latter have been obtained by running the TCP connections under the bandwidth partitions provided by the TEAM software, and confirm the good prediction properties of the TCP model used in the calculations. The values of the coefficients ρ and β (0.16 and 0.6, respectively) have been chosen with the criterion of maintaining a fairness factor always higher than 0.8 in the three cases. Under these conditions, it can be noted that the *Range*

strategy always exhibits a higher total goodput (anyway, this is not true in general).

As a final comparison, Fig. 2 show the goodput per connection in the three cases considered, for all classes, with the two selected values $\rho=0.16$ and $\beta=0.6$, and for all strategies. In Fig. 2 record #1 and #3, the goodputs corresponding to the *Tradeoff* strategy are always higher than those of the *BER Threshold* strategies. However, it is not so in Fig. 2 record #2, which makes it difficult to draw a general conclusion in this sense. In perspective, an adaptive choice of the coefficients, even within possible predefined limits (such as imposing a minimum threshold on the fairness, like we have done), may turn out to be advisable. This is the matter of current investigation. What can be further observed is that, in all cases considered, the range strategy tends to put a stronger penalization on stations suffering from heavy fading (as is the case with station # 4), and to more evenly favor the other stations.

5. 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 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. Further research is currently ongoing, to apply the proposed strategies in a dynamic environment, where fading levels and number of TCP connections in the system change over time in unpredictable fashion, as well as in traffic engineering for multiservice satellite networks.

REFERENCES

1. 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.
2. N. Celandroni, F. Potorti, "Maximising single connection TCP goodput by trading bandwidth for BER", *Internat. J. Commun. Syst.*, vol. 16, no. 1, pp. 63-79, Feb. 2003.
3. N. Celandroni, E. Ferro, F. Potorti, "Goodput optimisation of long-lived TCP connections in a faded satellite channel", *Proc. IEEE Vehic. Technol. Conf.*, Milan, Italy, May 2004.
4. 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.
5. 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.
6. H. J. Chao, X. Guo, *Quality of Service Control in High-Speed Networks*, John Wiley & Sons, New York, NY, 2002.
7. F. Alagoz, D. Walters, A. Alrustamani, B. Vojcic, R. Pikholtz, "Adaptive rate control and QoS provisioning in direct broadcast satellite Networks," *Wireless Networks*, vol. 7, no. 3, pp. 269-281, 2001.
8. 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.
9. The Network Simulator – ns2. Documentation and source code from the home page: <http://www.isi.edu/nsnam/ns>.
10. 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.
11. <http://wnet.isti.cnr.it/software/tgep.html>.
12. "Q1401: K=7 rate 1/2 single-chip Viterbi decoder technical data sheet", Qualcomm, Inc., Sept. 1987.
13. J. G. Proakis, "Digital communications", McGraw Hill, New York, NY, 1995.
14. K. W. Ross, *Multiservice Loss Models for Broadband Telecommunication Networks*, Springer, London, UK, 1995.
15. 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.
16. <http://www.eutelsat.com/satellites/13ehb6.html>