A Study of Fairness Issues in the Assignment of Bandwidth to Elephant TCP Connections Sharing a Rain-Faded GEO Satellite Channel

Nedo Celandroni*, Franco Davoli°, Erina Ferro*, Alberto Gotta°
*ISTI-CNR, Area della Ricerca del C.N.R., Via Moruzzi 1, I-56124 Pisa, Italy
nedo.celandroni@isti.cnr.it, erina.ferro@isti.cnr.it
°DIST-University of Genoa and CNIT (Italian National Consortium for Telecommunications), University of Genoa Research Unit, Via Opera Pia 13, 16145 Genova, Italy
franco@dist.unige.it, alberto.gotta@isti.cnr.it

***Abstract*- The goodput of long-lived TCP connections can be maximized in error-prone channels by trading bandwidth for Bit Error Rate (BER). This fact can be exploited in the bandwidth and transmission parameters' assignment, by means of Adaptive Coding and Modulation (ACM), in satellite systems; in such a case, multiple source-destination (SD) pairs actually "see" channels with diverse characteristics, owing to unequal weather conditions. However, when dealing with multiple SD assignments, both aspects of goodput-maximization and fairness must be dealt with. Improving the effectiveness of TCP may be relatively easy, but maintaining the effectiveness while maintaining the fairness is a challenge. The paper compares alternative solutions, by taking into account different approaches.**

I. INTRODUCTION

Recently, it has been demonstrated that in error-prone channels the goodput of TCP connections can indeed be improved by trading bandwidth for BER [1]-[4], even without making any change to the TCP congestion control mechanism (at least, as regards long-lived TCP connections). Actually, having fixed a set of transmission characteristics (radio spectrum, antenna size, and transmission power), the selection of an appropriate modulation scheme and a FEC (Forward Error Correction) type allows choosing the BER (Bit Error Rate) and the IBR (Information Bit Rate) of the link that maximize the goodput of a long-lived TCP connection. This optimization can be done for different channel quality conditions caused, for example, by atmospheric events. Indeed, over wireless links with a given bandwidth, any gain in BER (and, consequently, in the packet loss) is generally obtained at expenses of the IBR, and the TCP goodput increases with the latter and decreases with the former. The FEC techniques adopted do not interfere in any way with the normal TCP behavior, as they are applied just before the transmission over the satellite channel. This fact was exploited further [5]-[7], in the case of long-lived (elephant) TCP connections crossing a satellite system, where different SD pairs actually see channels with different characteristics, owing to weather conditions in diverse geographical areas. Essentially, ACM is applied as a specific fade countermeasure at the physical layer (see, e.g., [8] for a short account on satellite fade countermeasures in general), consisting of bit and code rate adaptation, with the awareness that long-lived TCP connections are being served, and adjusted accordingly. In the absence of information on the kind of traffic being handled, the aim of ACM is to keep the BER below a given level (if possible), independently of the bandwidth available (i.e., satisfying a pure physical layer constraint). As regards the possibility of segregating long-lived TCP connections from other traffic, as well as from short-lived ones, a look to references [9]-[11] can help.

However, when assigning bandwidth and ACM parameters to different SD pairs, with the aim to obtain a performance optimization of the whole system, two conflicting goals arise, namely, TCP goodput maximization and fairness in the allocation. This issue has been considered in some detail in [7], where different techniques are defined and compared, with respect to these two criteria. However, a systematic study on the impact of the above described resource allocation on fairness was not attempted, even in the static case, where fading conditions at the earth stations are assumed to remain stationary for a long time. The aim of the present paper is to perform a deeper investigation on this topic.

The paper is organized as follows. In the next Section, we shortly describe the fading and TCP models adopted. The bandwidth allocation methods and some fairness definitions are recalled in Sections III and IV, respectively. Section V presents the results of our numerical study and section VI concludes the paper.

II. SATELLITE FADING AND TCP MODEL

In our fully meshed satellite network, where N fixed earth stations use the Ka band of a geostationary satellite transponder as a bent-pipe channel, fade attenuations of the signal, due to bad weather conditions, are counteracted by applying adaptive FEC (Forward Error Correction) codes and modulation rates to the data before their transmission on the satellite channel, according to the detected attenuation level of the signal. This induces a redundancy of the data at the physical level. This solution, which is widely adopted, has an impact on the bandwidth usage, which affects the bandwidth allocation policies at the other layers. Data transmitted may have a real-time nature or they may be inherent to TCP connections. When transmitted over a faded satellite channel, data need to be protected according to the fade level of the receiving station (as regards the uplink, we assume to operate a power control). In order to avoid too many oscillations in applying the fade countermeasures according to each single fade level variation (which can be in the order of a second) of the receiving stations, the transmitting station categorizes each values of the signal attenuations *f, as measured and communicated by the receiving stations,* in *F* classes (*f*=1,…, *F*). The countermeasure strategy adopted remains unchanged for all those levels of signal attenuation (or, equivalently, of carrier power to one-sided noise spectral density ratio C / N_0) that belong to the same class. Thus, for each type of traffic, a *fade class* aggregates those fade levels that need the same data redundancy to sustain the QoS required by the relevant application. The redundancy (possibly dependent on the type of traffic being served) is expressed at station i by coefficients $r_f^{(i)}$, *f*=1,…, *F*, which represent the ratio between the IBR in clear sky and in the specific working condition, respectively. From now on, we will assume the redundancy to be applied before the transmission of IP packets that leave buffers dedicated to serve the TCP elephants of each specific SD pair, which have been identified and segregated from other traffic types.

The complete description and formulas of a model that can be used to derive analytical expressions of the normalized (to the bottleneck rate) *goodput* T_g of long-lived TCP connections over a satellite link can be found in [5]-[7], while the classical Additive White Gaussian Noise (AWGN) approach is used to model the satellite channel. We denote the absolute goodput by $\hat{T}_g$, i.e.

$$\hat{T}_g = T_g\left(\mu/n\right) = T_g\left(1/n\right)\cdot\left(B/r_f\right) \tag{1}$$

where B is the link rate in segments/s in clear sky conditions, μ [segments/s] is the bottleneck rate on the satellite link, and n is the number of TCP sources, which experience the same delay and get an equal share of the link [12]. The goodput expression depends, among other parameters, on the link bandwidth B and on the segment loss rate q (which is a function of the BER and, hence, of the transmission parameters)

III. The allocation methods

We refer to connections inherent to the same SD pair, which experience a specific channel condition, as belonging to the same "class"; as already mentioned, they feed a common buffer at the IP packet level in the earth 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 or from other earth 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 condition. This combination gives rise to the corresponding redundancy coefficient for those connections, which appear in relation (1).

The goal of the allocation is to satisfy a global optimality criterion, which involves *goodput* and *fairness* among the connections. Therefore, in correspondence of a specific channel situation, and a given traffic load, we face a two-criteria optimization problem. The decision variables are the service rates of the above-mentioned IP buffers for each SD pair, and the corresponding transmission parameters.

In [5]-[7] we developed and studied a group of allocation strategies, referred to as CLARA (*C*ross *L*ayer *A*pproach for *R*esource *A*llocations), which includes four approaches (named *range, tradeoff, BER threshold,* and *merge*), whose performance in terms of goodput and fairness in the allocation were compared. The approach named "*range*" performed better in all situations considered; thus, in this paper we refer to it to compare the fairness properties of different allocations. *Range* allocates the bandwidth by first computing the fair allocation (and the corresponding redundancy), i.e., the goodput-equalizing one; then, it chooses the bandwidth-redundancy pairs that maximize the goodput, within a given range of bandwidth values around the previous allocation. In [6], [7] this approach has been compared with another methodology, which considers the maximization of the sum of the logarithms of the connections' goodputs (optimized with respect to the redundancy for each value of the bandwidth allocation). The latter corresponds to seeking a Nash Bargaining Solution over the connections, which has intrinsic fairness properties; the resulting allocation has been termed *Generalized Proportionally Fair* (GPF), following [13]. The different strategies of the CLARA group, including the GPF, have been compared in a static fading scenario first, and then in a dynamically varying one, with fading traces taken from real-life samples [7].

IV. Fairness definitions

In order to compare different choices in terms of fairness, in [5]-[7] we defined and used the following index:

$$\varphi_f = 1 - \frac{\sum_{j=1}^{L}\left|\hat{T}_g^{(j)} - \overline{T}_g\right|}{2\overline{T}_g(L-1)} \tag{2}$$

where $L=\sum_{i=1}^{F} n_c^{(i)}$ is the total number of ongoing TCP connections ($n_c^{(i)}$ being the number of connections of class *i*), and $\overline{T}_g=\frac{1}{L}\sum_{k=1}^{L}\hat{T}_g^{(k)}$ is the average goodput.

We recall here two other fairness concepts: i) the proportionally fair one (see [14], [15] and generalizations in [16], [17]), and the fairness index by Chiu and Jain [18]. An assignment $B_1^*,...,B_F^*$ (in terms of bandwidth) is said to be *Weighted Proportionally Fair* with weight vector $\left(w_1^*,...,w_F^*\right)$ if

$$\sum_{j=1}^{F} w_j^* \frac{B_j - B_j^*}{B_j^*} \leq 0 \tag{3}$$

where $B_1,...,B_F$ is any other assignment; the assignment is said to be *Generalized Proportionally Fair* with respect to utility functions $f^{(j)}$ [13] if:

$$\sum_{j=1}^{F} \frac{f^{(j)}\left(B_j\right) - f^{(j)}\left(B_j^*\right)}{f^{(j)}\left(B_j^*\right)} \leq 0 \tag{4}$$

In particular, we have chosen in [6], [7]

$$f^{(j)}\left(B_j\right) = \left[\hat{T}^{(j)^\circ}\left(B_j\right)\right]^{n_c^{(j)}} \tag{5}$$

where $\hat{T}^{(j)^\circ}\left(B_j\right)$ is the maximum of the goodput with respect to the redundancy assignment for the given bandwidth B_j.
The fairness index of Chiu and Jain is defined as

$$\varphi_{f,CJ} = \frac{\left(\sum_{j=1}^{L} \hat{T}_g^{(j)}\right)^2}{L \sum_{j=1}^{L} \left(\hat{T}_g^{(j)}\right)^2} \tag{6}$$

In the present paper, we adopt the index in (6), which is a "classical" one and has been extensively used, to compare the *GPF* and *range* strategies in terms of fairness and flexibility.

V. Numerical Results in the Stationary Fading Case

We consider the cross-layer optimization for the achievement of the "best" compromise between the maximization of TCP goodput and fairness. The numerical details are the same as in [7], with a combination of long-lived TCP Reno connections, sharing various bottleneck links, determined by 10 different fading classes, under the HB6 satellite link budget [19]. However, the results we discuss here are completely new, as they are relevant to a more thorough analysis of the stationary fading case, and to the application of the fairness index defined in (6).
The TEAM (TCP Elephant bandwidth Allocation Method) software [20] has been used in the calculations. It implements a number of allocation methods, by performing the maximization of the corresponding performance indexes (that involve functions of the goodput or of goodput and fairness) over a discrete set of bandwidth allocations and transmission parameters. In particular, the *range* allocation procedure involves the following operations:
1. Compute the pairs $\left(\overline{B}_i\ ,\ \overline{r}_i\right)$, $i=1,...,F$, corresponding to the goodput-equalizing fair choice;
2. Choose a "range coefficient" $\beta \geq 0$;
3. Compute the global allocation, by effecting the constrained maximization of the sum of the goodputs (under the linear constraint that the bandwidth allocations to the classes sum to the total satellite bandwidth *W*), with $\overline{B}_i$ varying in the range $\left[\max\left(\overline{B}_i(1-\beta),\ 0\right),\ \min\left(\overline{B}_i(1+\beta), W\right)\right]$.
On the other hand, as already mentioned, the *GPF* allocation performs the maximization of the sum of the logarithms of functions (5). The granularity in the discretization of the bandwidth is defined in terms of the *minimum bandwidth unit* (*mbu*) that can be allocated. In the present case, this corresponds to 5 kbit/s, and the total satellite bandwidth is 1600 *mbu*.
The comparison of the two techniques has been performed in a number of 23 "trials", each involving 10 classes of SD pairs. The connections of each class experience different values of C/N_0, each of which equals the preceding one, plus 1 dBHz, in increasing order from class 1 to class 10, starting with C/N_0=62.0 dBHz for class 1 and 71 dBHz for class 10 in the first trial. All values of C/N_0 increase by 0.5 dBHz with each trial, from 1 to 23. In practice, we start from a situation that sees the majority of classes in deep fade, and end up with the majority of classes in clear sky. The comparisons are effected under different numbers of connections per class (all classes have the same number of connections).

The fairness achieved (automatically) by the GPF allocation is assumed as a target, and is computed according to (6); the β coefficient of the *range* procedure is then chosen so that the fairness (always in the sense of (6)) of its allocation be never below the target value (if possible).

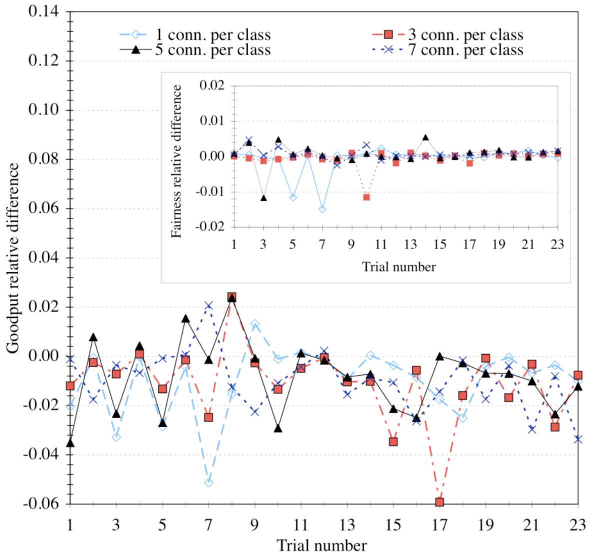

Fig. 1. Relative goodput difference of *range* with respect to *GPF*, when the former targets the fairness of the latter; small deviations from the target are shown in the inner graph.

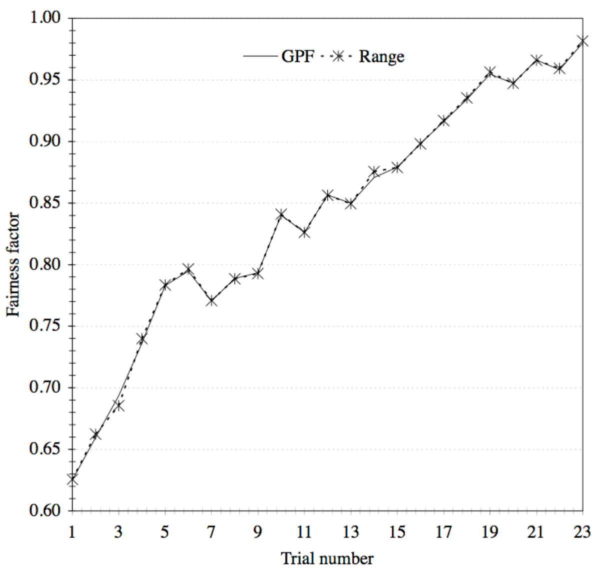

Fig. 2. Fairness factors of the *GPF* and *range* allocations, respectively.

Fig. 1 shows the normalized (to *GPF*) difference between the overall goodput achieved by the *range* allocation and the one achieved by the *GPF* allocation, for 4 different numbers of connections per class, under the procedure described above to maintain comparable fairness levels between the two allocations. It is evident that the global goodputs achieved by the two strategies differ for a few percent, under all fading conditions. Fig. 2 shows the almost perfect correspondence in fairness index between the two allocations.

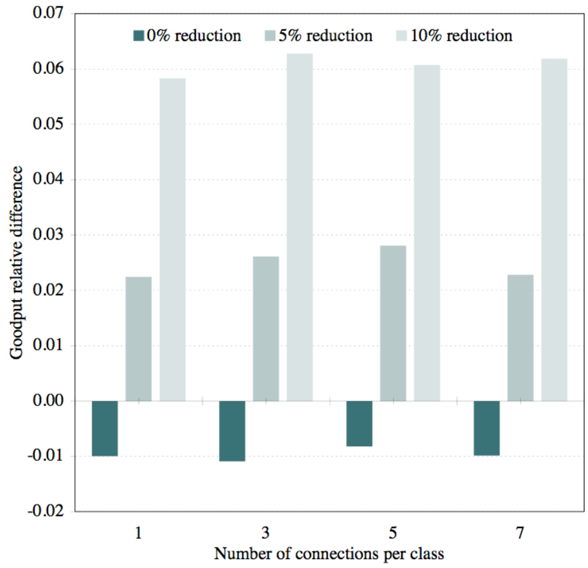

Fig. 3. Relative goodput difference between *range* and *GPF* allocations, averaged over all trials, by accepting a fairness reduction of *range* with respect to *GPF*.

On the other hand, an advantage of *range* is the possibility of controlling the tradeoff between goodput and fairness, by means of the β coefficient, whereas no tradeoff is possible under the *GPF* allocation.

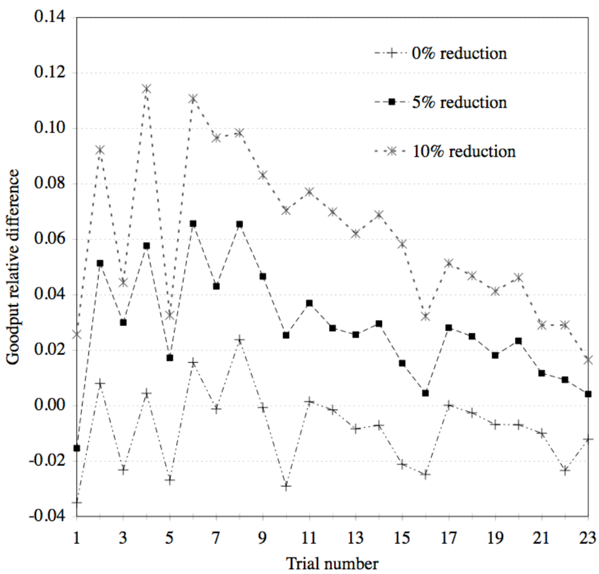

Fig. 4. Relative goodput difference between *range* and *GPF* allocations, by accepting a fairness reduction of *range* with respect to *GPF*.

This fact is highlighted in Fig. 3, where the global relative goodput difference (over all trials) is represented, by accepting a given reduction in fairness for the *range* allocation, with respect to the *GPF* one.

In this case, whereas the *GPF* shows a small advantage at the same fairness (0% reduction), there is a gain in goodput of the *range* allocation up to about 6%, by renouncing to some degree of fairness (10% reduction).
The same situation is depicted separately for each trial in Fig. 4, assuming 5 connections per class.
As a final remark, it is worth noting that, in performing the bandwidth allocation, the achievement of a minimum bandwidth efficiency per connection (ratio of the connection's goodput to its available bandwidth) is checked. This deserves some further comments. In particularly extreme conditions of fading (low values of C / N_0), for the *range* assignment it may not be worth trying to keep the balance in the connections' goodput also for the heavily faded SD pairs, since this might imply a very low utilization of the bandwidth assigned. Therefore, we have decided to adopt an alternative solution, which consists of enforcing, for such SD pairs, the bandwidth that is necessary for their connections to achieve 50% utilization. This bandwidth can be computed by using the inverted relation of the goodput expression. The latter (see [5]-[7]) is actually given by two different formulas, according to the channel parameter values (round trip time, segment loss rate, bottleneck rate): one corresponding essentially to the goodput expression derived in [21], the other to an approximating polynomial; for an efficiency per connection around 50%, the former can be applied, which is more easily inverted. After this assignment has been enforced (unless the bandwidth resulting for a given class is below 1 *mbu*, in which case the class is considered in outage), the *range* strategy is applied again to the remaining classes. This procedure is adopted for the initialization step of the *range* scheme, i.e., for the calculation of the values $\left(\overline{B}_i\ ,\ \overline{r}_i\right)$, $i = 1,...,F$, corresponding to maximum (close to 1) fairness. Then, whenever the coefficient β is enlarged and the allocation is sought that maximizes the global goodput within the given range, all classes are included again in the computation, since at this point their utilization can only increase.

VI. CONCLUSIONS

We have analyzed numerically the fairness behavior of two bandwidth allocation methods to TCP connections, experiencing diverse fading conditions. In order to operate under similar fairness values, the fairness achieved by *GPF* is assumed as a target for the allocation with the *range* method. The main advantage of the latter consists of the possibility of independently tuning the bandwidth assignment (and the corresponding transmission parameters), by means of the choice of a coefficient. In other words, the method allows trading goodput at the price of slightly unbalancing the fairness. The analysis has been conducted over a wide range of values of the carrier-to-noise spectral density ratio, representing the channel condition, and the flexibility of the *range* method has been highlighted.

ACKNOWLEDGMENT

This work was partially supported by the Italian Ministry of Education, University and Research (MIUR) and the National Research Council (CNR), in the framework of the "*IS-MANET*" project, and by the European Commission in the framework of the NoE project "*SatNEx*" (contract no. 507052).

REFERENCES

[1] 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.

[2] C. Barakat, E. Altman. "Bandwidth tradeoff between TCP and link-level FEC", *Computer Networks*, vol. 39, no. 5, June 2002.

[3] D. Barman, I. Matta, E. Altman, R. El Azouzi, "TCP optimization through FEC, ARQ, and transmission power tradeoffs", *Proc. 2nd Internat. Conf. on Wired-Wireless Internet Commun. (WWIC'04)*, Frankfurt (Oder), Germany, Feb. 2004.

[4] L. Galluccio, G. Morabito, S. Palazzo, "An analytical study of a tradeoff between transmission power and FEC for TCP optimization in wireless networks", *Proc. IEEE INFOCOM 03*, San Francisco. CA, March 2003, vol. 3, pp. 1765-1773.

[5] 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", in D. Gaïti, S. Galmés, R. Puigjaner, Eds., *Network Control and Engineering for QoS, Security and Mobility, III*, Springer, New York, NY, 2004, pp. 83-96.

[6] N. Celandroni, F. Davoli, E. Ferro, A. Gotta, "Adaptive bandwidth partitioning among TCP elephant connections over multiple rain-faded satellite channels", *3rd Internat. Workshop on QoS in Multiservice IP Networks*, Catania, Italy, Feb. 2005; in M. Ajmone Marsan, G. Bianchi, M. Listanti, M. Meo, Eds., *Lecture Notes in Computer Science*, 3375, Springer, Berlin, 2005, pp. 559-573.

[7] N. Celandroni, F. Davoli, E. Ferro, A. Gotta, "Long-lived TCP connections via satellite: cross-layer bandwidth allocation, pricing and adaptive control", *IEEE/ACM Trans. Networking* (to appear).

[8] L. Castanet, A. Bolea-Alamanac, M. Bousquet, "Interference and fade mitigation techniques for Ka and

Q/V band satellite communication systems", *Proc. Internat. Workshop of COST Actions 272 and 280 on Satellite Communications – From Fade Mitigation to Service Provision*, ESTEC, Noordwijk, The Netherlands, May 2003 [available online: http://www.cost280.rl.ac.uk/documents/WS2%20Proceedings/documents/pm-5-002.pdf].

[9] S. Ebrahimi-Taghizadeh, A. Helmy, S. Gupta, "TCP vs. TCP: A systematic study of adverse impact of short-lived TCP flows on long-lived TCP flows", *Proc. IEEE INFOCOM '05*, Miami, FL, March 2005.

[10] S. Yilmaz, I. Matta, "On class-based isolation of UDP, short-lived and long-lived TCP flows", Technical Report BU-CS-2001-011, Boston University, Computer Science Department, June 2001.

[11] K. Papagiannaki, N. Taft, C. Diot, "Impact of flow dynamics on traffic engineering design principles", *Proc. IEEE INFOCOM '04*, Hong-Kong, March 2004, vol. 4, pp. 2295-2306.

[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] C. Touati, E. Altman, J. Galtier, "On fairness in bandwidth allocation", Research Report no. 4296, INRIA, Sophia Antipolis, France, Sept. 2001.

[14] F. P. Kelly, "Charging and rate control for elastic traffic", *Europ. Trans. Telecommun.*, vol. 8, pp. 33-37, 1998.

[15] F. P. Kelly, A. Maulloo, D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability", *J. Op. Res. Society*, vol. 49, pp. 237-252, 1998.

[16] J. Mo, J. Walrand, "Fair end-to-end window based congestion control", *IEEE/ACM Trans. Networking*, vol. 8, no. 5, pp. 556-567, Oct. 2000.

[17] R. J. La, V. Anantharam, "Utility-based rate control in the Internet for elastic traffic", *IEEE/ACM Trans. Networking*, vol. 10, no. 2, pp. 272-286, April 2002.

[18] D. Chiu, R. Jain, "Analysis of the increase and decrease algorithms for congestion avoidance in computer networks", *Computer Networks and ISDN Syst.*, vol. 17, pp. 1-14, 1989.

[19] http://www.eutelsat.com/satellites/13ehb6.html

[20] http://www.isti.cnr.it/ResearchUnits/Labs/wn-lab/software-tools.html

[21] 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, no. 2, pp. 133-145, April 2000.