

Networking with Multi-Service GEO Satellites: Cross-Layer Approaches for Bandwidth Allocation[♦]

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- Cross-layer Radio Resource Management (CL-RRM) has been recently investigated by quite a few research groups in wireless communications. In the specific satellite networking environment, the paper presents an overview of different CL-RRM techniques devoted to dynamic bandwidth allocation, whose interactions span the physical, data link, network and transport layers, in various combinations. A multi-service setting is considered, in the presence of variations in both traffic and channel conditions. Regarding the latter, bit and coding rate adaptation are adopted as fade countermeasure, and their effect on the higher layers is modeled as a bandwidth reduction. Traffic models and methodologies for dynamic bandwidth allocation and performance optimization are discussed. Numerical examples are presented to highlight throughput/fairness tradeoffs for long-lived TCP connections that share multiple channels with different fading depth.

I. INTRODUCTION

Radio Resource Management (RRM) is emerging as an essential component of modern wireless networking systems, whose general goal is the optimization of all available resources (energy, bandwidth, storage and computational power) in the provision of Quality of Service (QoS) to different information sources. This applies to Wireless LANs [1] and cellular systems [2, 3], as well as to satellite networks [4, 5]. In particular, when satellite systems are involved, the network not only has to face variable load

[♦] Work supported by the European Commission in the framework of the SatNEX Network of Excellence (contract no. 507052). This paper is an extended version of reference [16].

multimedia traffic, but also variable channel conditions and large propagation delays. The variability in the characteristics of satellite links is mainly caused by the variable signal attenuation, due to bad atmospheric events, which particularly affect transmissions in the Ka band (20-30 GHz). It is therefore crucial to make use of adaptive network management and control algorithms to maintain the QoS of the transmitted data.

On the other hand, as regards the presence of multiple services within the network, two main aspects need to be considered: sources with different statistical behaviour, and a multiplicity of performance requirements. In general, especially in the QoS-IP world, one of the main points associated with multiple services is the identification of some types of flows, with specific statistical and performance characteristics. A very basic distinction is between *inelastic* and *elastic* traffic. Inelastic traffic is generally time-sensitive, and therefore it needs some bandwidth guarantees; we will refer to it as Reserved Bandwidth (RB) traffic. Among RB flows, a further distinction can be made between those requiring minimum delay/jitter and exhibiting a Constant Bit Rate (CBR) behaviour (e.g., voice with no silence detection or CBR-encoded video), and the Variable Bit Rate (VBR) ones, which are bursty in nature, though still requiring the satisfaction of real-time constraints. Elastic traffic is a part of the Best-Effort (BE) category; long-lived TCP connections (also called “elephant”) are of this type, where the “elastic” behaviour is induced by the TCP congestion control mechanism. Short-lived TCP connections still belong to BE, but they do not present an elastic behaviour, due to their short duration, which most of the times prevents exiting the slow-start phase (inelastic BE traffic). In addition, inelastic BE traffic may also include UDP flows with no particular bandwidth guarantees (non-real-time VBR). The presence of multiple traffic types in the network increases its heterogeneity and, as a consequence, complicates the task of network control, which is exerted to distribute the resources (in particular, the bandwidth), in order to achieve some desired goals. Actually, QoS in a strict sense tends to be currently perceived as a less critical issue in cabled networks than in the past (see, e.g., [6, 7]); however, in the wireless environment, the effects of dynamic channel variations add to those of varying traffic patterns, and, together with the limited bandwidth, make the enforcement of QoS requirements much more challenging.

A further aspect that has recently stimulated much interest in the wireless scenario is related to cross-layer optimization. The combined action among various layers of the network is likely to improve the end-to-end performance, through a coordination of control actions, QoS mapping, and cross-layer information exchange for control purposes. Designing the layers' architecture, possibly with the adoption of ad-hoc solutions, becomes of paramount importance in a heterogeneous network, which may sometimes require the separation of the satellite component (e.g., by introducing *performance-enhancing proxies* or *relay entities*), in order to work effectively. Much care has to be taken in these approaches, in order not to disrupt wise layering principles, possibly leading to unstable system behaviour [8]. Although more rarely in a cross-layer context, the dynamic control of resource allocation has been widely considered in the literature (e.g., [9-26] in the satellite environment), as regards the allocation of bandwidth, transmission power, antenna beams, or combinations thereof.

This paper summarizes some specific dynamic bandwidth allocation problems over a geostationary (GEO) satellite network loaded with multi-service IP traffic. The main objective is to describe some classes of parametric optimization problems in this context, as well as possible solution techniques, in a unified setting, and to highlight the cross-layer approach. The operative environment we refer to is a Multi-Frequency Time Division Multiple Access (MF-TDMA) system, where we assume that a station cannot transmit at different frequencies in the same temporal slot. In particular, we consider cross-layer interactions between the physical and the Medium Access Control (MAC) layers, between the network and the MAC layers (QoS mapping), and finally between the physical and the transport layers.

The paper is organized as follows. The next Section deals with the physical layer, and introduces the fade countermeasure technique adopted, along with its effect on higher layers. In Section III we describe some traffic models. Sections IV and V deal with specific CL-RRM (Cross-Layer Radio Resource Management) techniques, first in the case of mixed RB and BE traffic, and then in the case of long-lived TCP connections; in both cases, diverse fading conditions are experienced. Section VI presents an instance of numerical results in the TCP case, while Section VII concludes the paper.

II. FADE COUNTERMEASURES BY ADAPTIVE CODING AND THE BANDWIDTH REDUCTION EFFECT

We consider a fully meshed satellite network, where N fixed earth stations use the Ka band of a GEO satellite transponder as a bent-pipe channel. Fade attenuations of the signals, due to bad weather conditions, are countered by applying adaptive FEC (Forward Error Correction) codes and bit rates, i.e., data are redounded at the physical layer before their transmission on the satellite channel, according to the detected attenuation level of the signal. This is just one of many possible actions that can be taken in this respect (see, e.g. [27]). We focus on this type of fade countermeasure because it has an impact on the bandwidth usage that affects the bandwidth allocation policies at the other layers; this is why this type of fade countermeasure has been exploited in some works in CL-RRM. When all stations address their traffic to the same destination station, such as in a VSAT system, the transmission parameters only depend on the data source. The same applies when the power budget is of “up-link predominant” type, i.e., it is insensible to down-link fading. More generally, destination stations may experience different down-link attenuations, and the link budget may be more balanced. Transmitting stations may thus apply different FEC codes to the data and transmit data at different bit rates, according to the down-link fades of the data destinations; in this case, we refer to source-destination (SD) pairs.

In general, the channel condition is expressed by the carrier power to one-sided noise spectral density ratio C/N_0 . If adaptive coding is applied, the redundancy added to the data for each particular value of C/N_0 may depend on the specific application the data refer to (at least, when we are operating in a cross-layer fashion). We distinguish two specific cases of data transmission: i) data deriving from RB and BE inelastic traffic; ii) data deriving from BE traffic originated by long-lived TCP applications with elastic behaviour. Obviously, a condition for that is the possibility to distinguish the different flows, as will be mentioned very briefly in the next Section.

We consider case i) first. For this type of traffic, the usual QoS requirement at the physical layer is to keep the Bit Error Rate (BER) below a given threshold. In the presence of fading variations (i.e., variations in C/N_0), this requirement can be satisfied by means of adaptive coding. The signal fade may vary in very short time

intervals, even less than a second. In order to avoid too many oscillations in applying the fade countermeasure, according to each single fade level variation, the measured value of the signal attenuation is categorized in a class $f, f=1, \dots, F$, where F is the number of fade classes. The countermeasure strategy adopted remains unchanged for all those levels of signal attenuation (or, equivalently, of 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, obtained in general by a combination of bit transmission rate and coding rate, is expressed at station i by coefficients $r_f^{(i)}, f=1, \dots, F$, which represent the ratio between the Information Bit Rate (IBR) in clear sky and the IBR in the specific working condition. As BE and RB traffic have different QoS requirements, we indicate the respective redundancies with $r_{f,be}^{(i)}$ and $r_{f,rb}^{(i)}$. It is worth noting that the inverses of the redundancy coefficients can be seen as multiplicative factors that express the reduction in net bandwidth caused by the countermeasure adopted.

As regards case ii), i.e., the long lived TCP connections, we will see in Section V that a tradeoff exists between bandwidth and BER in the maximization of the TCP end-to-end transfer rate (goodput). Therefore, in this case it is no longer true that a given redundancy corresponds to a set of C/N_0 values. We will therefore use the term “class” in a different sense, just to indicate a group of TCP connections that experience the same C/N_0 . Nonetheless, we can still use the redundancy coefficients to represent given combinations of bit and coding rates, whose application with a certain bandwidth and C/N_0 would yield a corresponding BER.

The effect on the bandwidth of the adaptive coding fade countermeasure technique that we summarized should be taken into account in traffic engineering of satellite networks. For MF-TDMA systems, as the ones we are considering, the redundancy coefficients can be computed from the link budget parameters [19], by deriving the channel bit energy to one-sided noise spectral density ratio (E_c/N_0) from C/N_0 , and thus deriving the corresponding BER over all C/N_0 values of interest between the extreme situations of clear sky and outage. The outage condition may be defined in different terms, according to the specific application considered: BER above a given threshold for

inelastic traffic, or goodput (i.e., the end to end transfer rate of long-lived TCP connections) below a minimum acceptable value.

III. TRAFFIC MODELS AND PERFORMANCE INDEXES

A. Mixed (RB/BE) traffic.

The RB traffic is characterized by bandwidth requests for flows that need some priority at the satellite link level, within some specific DVB (Digital Video Broadcasting) service classes [28, 29]. The flows can be voice or MPEG video, bandwidth reserved for DiffServ aggregates, Multi Protocol Label Switching (MPLS) pipes, and the like. Whenever such flows are characterized by a relatively low burstiness (e.g., the Peak-to-Average ratio of their rates is close to 1, as is the case with CBR), DAMA (Demand Assignment Multiple Access) schemes can be efficiently used to schedule the transmission from the earth stations at the MAC layer [21]. In this case, bandwidth can be allocated by using Call Admission Control (CAC) functions. The dynamics of interest is at the connection level, and the relevant performance index is the steady-state *blocking probability* (P_{block}). For RB traffic we can adopt the usual birth-death model with exponential distribution of call inter-arrival times (Poisson arrivals). If all connections of station i have the same peak rate $B^{(i)}$, and also belong to the same fade class f , we face a particular single-class case, where the expression of the blocking probability is given by the classical Erlang B loss formula [30]. By using this formula, the maximum number of acceptable calls $N_{\max}^{(i)}$ can be easily derived at station i , given the Erlang traffic intensity and a desired upper bound on the blocking probability. In this particular case, under a specific fading condition that requires redundancy $r_{f,rb}^{(i)}$, the total bandwidth needed to carry the maximum number of calls is $B^{(i)} r_{f,rb}^{(i)} N_{\max}^{(i)}$.

In the case of multiple fade classes per station (and/or multiple CBR services with different peak rates), data transmission would require the simultaneous application of different redundancy values. In this multi-class case, the blocking probability results from the solution of a *stochastic knapsack* problem [30], or from some approximation of it. The situation can be always kept in the single class case if the bandwidth is assigned separately *per-fading-* or *per-traffic-class* inside the station.

Various traffic models have been used to represent the burst-level behaviour of real-time VBR traffic; among these, there are *voice with silence detection* and *VBR-encoded MPEG video*. In the presence of VBR traffic, two control functionalities at different time scales should be employed, namely, CAC and dynamic bandwidth allocation at the burst level, in order to guarantee at the same time both a specified degree of QoS and an efficient bandwidth utilization. In the next Section we will mention two techniques that have been proposed in this setting (also employed in a cross-layer context); they are *ARAM* (Adaptive Resource Allocation and Management) [11], and *CF/DAMA* (Combined Free/DAMA) [20, 21]. The burst-level model adopted in [11] for MPEG (1 and 2) video is based on the superposition of two first-order autoregressive processes with log-normally distributed noise sequences (2LAR) [31]. Since adjustments are possible in the service rate of admitted connections to face congestion events or channel variations, the QoS metric is the relative difference between the nominal rate and all rate reductions due to control or losses [11]. In [20], models capturing both Short Range Dependent (SDR) and Long Range Dependent (LRD) behaviours have been used to represent the arrival processes of traffic aggregates to the User Terminal queues in a DiffServ setting. They are based on Markov Modulated Poisson Processes (MMPP) and Pareto Modulated Poisson Processes (PMPP), giving rise to MMPP/G/1 and PMPP/G/1 queuing systems, respectively. The QoS metric adopted is the probability of the queue length of each service queue to exceed a given threshold, depending on the service; this probability must be kept below a specified value, beyond which the station is considered in outage. The scheduling of the MAC queues must be such that this constraint is maintained for the IP-level queues (i.e., those corresponding to Expedited Forwarding (EF), Assured Forwarding (AF), and BE services for a given earth station). The control process works upon requests for bandwidth allocation at the burst level; these requests can be satisfied within a Round Trip Time (RTT), the time necessary for the request to reach the scheduler and the response to be received (referred to as Dynamic Capacity Allocation -DCA- cycle time in [20]). If the state of the sources can be assumed changing more slowly than the DCA cycle time, within which the allocated bandwidth remains constant, the queuing behaviour in these intervals can be approximated by a much simpler M/D/1 system.

In general, IP data packets are fragmented at layer 2 into fixed-size cells (ATM or DVB [29]) before transmission on satellite. As regards inelastic BE traffic, it can be assumed that, at each station i , such cells are queued in a finite buffer of capacity $Q^{(i)}$, served with no particular guarantees. In this context, the quantity of interest is the *cell loss probability* (P_{loss}) in the queue of station i . An approximate evaluation of this quantity has been used in [18] and [19], under a discrete-time self-similar traffic model, derived from [32]. This model represents the superposition of on-off sources, whose active periods (bursts) have Pareto-distributed ‘on’ times. In the presence of both RB and inelastic BE traffic, as considered in [18] and [19], the extraction rate is determined by the residual capacity $C_{be}^{(i)}(t)$, out of the total capacity $C^{(i)}$ allocated to station i , available for the inelastic BE traffic after serving all RB connections in progress at the required transmission rate (as determined by the peak or by an “equivalent bandwidth” [30] figure, as well as by the current redundancy value). In general, the residual bandwidth is a random variable; as a consequence, the loss probability can be considered only as conditional on the number of RB connections in progress; its average must be computed with respect to the statistics of the Markov chain that describes the connection dynamics. Thus, the performance measure for inelastic BE traffic is the average P_{loss} . The cross-layer action on this quantity stems from the redundancy coefficients (in this case of both RB and inelastic BE traffic types), which parameterize the P_{loss} analytical expression. Even if redundancy coefficients are dynamically varying quantities, the analytical expressions of the performance indexes are based on stationary considerations, i.e., as if the instantaneous value of each redundancy would last forever. In reality, the performance indexes are re-computed at each change of fade class, which implies new redundancy values. Dynamic control schemes based on this model fall into the category of “*open-loop feedback*” repetitive control [33], where the initial time of the control horizon continuously shifts ahead.

B. Long-lived TCP traffic.

We sketch now the main structure of a model that can be used to derive analytical expressions of the goodput of long-lived TCP Reno/NewReno connections over a satellite link. Let μ [segments/s] be the bottleneck rate (the satellite link), n the number of TCP sources (assumed to experience the same delay and to get an equal share of the

link), $d = c_l + 1/\mu$ the round trip delay between the beginning of the transmission of a segment and the reception of the relative acknowledgement (c_l being twice the channel latency), and q the segment loss rate (assuming independent losses). Let also b and T_o be the number of segments acknowledged by each ACK received, and the timeout estimated by the sender TCP, respectively. The TCP connections that share the same link also share an IP buffer, of capacity at least equal to the product μd , inserted ahead of the satellite link. It is worth noting that these connections can be actually recognized and separated from short-lived ones [34]. It has been shown in [17] that the parameter $y = q(\mu/n)^2 d^5$ plays a major role in the expression of the TCP goodput. Indeed, for $y > 1$, the relative goodput (normalized to the bottleneck rate) can be expressed as

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

which is derived by a slight modification of a classical formula [35]. Instead, for $0 \leq y \leq 1$, a good approximation of the goodput is given by the 4th order polynomial

$$T_g = 0.995 + 0.11y - 1.88y^2 + 1.98y^3 - 0.63y^4,$$

whose coefficients have been calculated by simulation and numerical interpolation [17], for $b = 1$.

In essence, by combining the analytical result (for $y > 1$) and the numerical approximation (for $0 \leq y \leq 1$), it is possible to obtain reasonably accurate analytical expressions of the goodput of a single connection, over a wide range of values of the parameters.

Then, assuming to operate on an Additive White Gaussian Noise (AWGN) channel, the segment loss rate q can also be approximated as a function of the BER p_e , the segment length in bits l_s , and the average error burst length (ebl) l_e [36]; namely, $q = 1 - (1 - p_e/l_e)^{l_s}$. By interpolating numerical values from [36], all these quantities can be analytically expressed as functions of the coding rates considered and of E_c/N_0 (channel bit energy to one-sided noise spectral density ratio), whose dB value is $C/N_0 - 10 \log_{10} b_r$, for a given bit rate b_r [17]. For each channel condition C/N_0 ,

among all available channel bit and coding rates that give rise to the same total redundancy coefficient r_f , the combination that yields the minimum BER is selected; thus, the segment loss rate q results a decreasing function of r_f . The TCP goodput relative to the bottleneck rate is a decreasing function of q ; as a consequence, it is an increasing function of r_f . The absolute goodput of a TCP connection \hat{T}_g is obtained by multiplying the relative value by the bottleneck rate: $\hat{T}_g = T_g \cdot (\mu/n) = T_g \cdot (1/n) \cdot (W/r_f)$, where W is the link rate in segments/s in clear sky conditions. Finally, for each channel condition there is a value of r_f that maximizes the absolute goodput of a TCP connection.

IV. CL BANDWIDTH ALLOCATION AND CAC IN THE PRESENCE OF RB AND BE FLOWS

Whenever we deal with flows, especially RB ones, some form of CAC should be exerted on incoming bandwidth requests, possibly in conjunction with bandwidth allocation among users and services. Ross in [35] provides a classification and an extensive discussion about alternative CAC solutions. Quite often, a functional form is a-priori decided for CAC strategies that map the available state information to acceptance decisions, thus transforming a functional optimization problem into a parametric one. Then, the “best” values of the parameters can be sought, in order to minimize a given cost function (or maximize a performance index). Different forms of this approach have been taken, among others, in [11], [18], and [19], which are all based on a cross-layer optimization. In [11], the presence of both VBR MPEG connections and Available Bit Rate (ABR) data has been considered. Here, CAC is exerted with the goal of keeping the probability that the bandwidth dedicated to VBR exceed a given value below a predetermined threshold. A bandwidth expansion factor, whose value is adaptively adjusted on the basis of measurements, is used to account for statistical multiplexing effects in VBR traffic. FEC and MPEG coding rate adjustments are other corrective actions taken to cope with traffic and channel variations. The approach taken in [18] and [19] considers RB and BE traffic; however, no rate adjustment derived from application-level coding is assumed to be available for RB flows. Adaptive cross-layer bandwidth partitions per station are derived, based on the performance indexes P_{block}

and P_{loss} , defined in III.A. The control architecture exhibits a hierarchical structure, where CAC tasks are delegated to local controllers at the stations, and capacity partitions $C^{(i)}$, $i=1, \dots, N$, are determined adaptively by a Master Control Station (MCS). It is worth noting that, owing to the dynamic fade changes, the assigned bandwidth $C^{(i)}$ may be temporarily insufficient to carry on the currently ongoing number of CBR connections $n^{(i)}(t)$ in the station (i.e., $C^{(i)} < B^{(i)}_{r_f, r_b}(t)n^{(i)}(t)$); since *inelastic* traffic is considered, in such cases one or more ongoing calls would be dropped. However, reallocations of the bandwidth partitions upon detection of significant changes in traffic intensities and fading classes do help in reducing the probability of this event.

As regards the MCS, the bandwidth allocation is formulated as an optimization problem in a discrete setting, with the assignment's granularity determined by the minimum bandwidth unit – *mbu*. If the performance index is a separable function of the stations' parameters (e.g., a sum of terms, each one depending only on the bandwidth to be assigned to a station), the problem can be solved numerically by applying Dynamic Programming over the stations [18], [19], possibly in a form that may greatly reduce the search space, by exploiting the presence of constraints.

Finally, it is worth noting that these model-based approaches can be by-passed by using a fluid approximation and treating the bandwidth partitions as continuous variables. A gradient descent technique can be adopted, in conjunction with Infinitesimal Perturbation Analysis (IPA) for gradient estimation [26]. The advantage of these methodologies is that they are measurement-based and require neither the knowledge of any functional form of the performance index nor any characterization of the traffic sources.

A CL-RRM problem involving the network and the MAC layers has been extensively considered in [20, 21]. In particular, DCA is applied by computing bandwidth requests for each earth station's DiffServ queues, which are passed to a centralized scheduler, typically residing in an MCS. The latter assigns the bandwidth proportionally to the requests received. The requests are computed on the basis of the queuing models and the performance indexes described earlier in III.A. The remaining capacity is assigned on a free basis, according to CF/DAMA. Only traffic and no fading variations are taken

into account, but, as is noted in [20], the effect of fade countermeasures might be included as a reduction in the available uplink bandwidth.

In a similar context, the problem of QoS mapping between adjacent layers has been recently treated in [37]. Rather than considering specifically the network and the MAC layers, the problem is posed in a more general setting, as defined by the ETSI BSM (Broadband Satellite Multimedia) protocol architecture [38, 39], at the SI-SAP (*Satellite Independent – Service Access Point*). Specifically, interworking between the Satellite Independent (SI) and Satellite Dependent (SD) architectural components is considered, by taking into account both the change in encapsulation format and the traffic aggregation (in the passage from SI to SD queues). In the presence of IP DiffServ queues at the higher layer, the problem consists in dynamically assigning the bandwidth (service rate) to each SD queue, so that the performance required in the SI IP-based SLA (Service Level Agreement) is guaranteed. By considering a fluid model and the loss volume as the performance indicator of interest, the IPA technique of Cassandras *et al.* [40] (already mentioned above in a different scenario) is applied, in order to maintain on-line the equalization between the loss volumes at the two different layers (by assuming that the resource allocation of SI is capable of satisfying the requirements). In doing so, both traffic and fading variations are taken into account. Also in this case, the effect of the latter appears through the bandwidth reduction coefficients defined in Section II.

V. CL BANDWIDTH ALLOCATION FOR TCP FLOWS

The application of adaptive FEC techniques has been investigated in [17], with the aim of optimizing the efficiency of TCP connections when transmitted over rain-faded geostationary satellite channels, with fixed user antennas. In this case, rather than reacting to fading changes by trying to keep the BER below a given threshold (BER Threshold methodology, THR), it is better to apply the same philosophy as in [36]. This consists in trading the bandwidth of the satellite link for the packet loss rate due to data corruption, in order to maximize the goodput of TCP connections. 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. The optimal transmission

parameters, for each channel condition, can be reported in look-up tables and then applied in an adaptive fashion. If connections take place between different SD pairs over satellite links, they may generally suffer from diverse fading conditions, according to the atmospheric effects on the source up-link and destination down-link. We refer to connections on the same SD pair, which experience a specific channel state, as belonging to the same “fade class”. At the IP packet level, these connections feed, a common buffer in the station, to which a transmission channel with specific characteristics corresponds. These characteristics may differ, in general, from those of other SD pairs originating either from the same or from other stations. Thus, in a station as many buffers are present as the station’s fade classes. 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 state. This combination gives rise to the corresponding r_f coefficient for those connections that appear in the expression of the goodput, according to the model mentioned in III.B.

The goal of the allocation algorithm is to satisfy a global optimality criterion, which involves goodput and fairness among the connections (where maximum fairness corresponds to dividing the bandwidth in such a way that all TCP connections achieve the same goodput). Therefore, in correspondence of a specific channel situation determined by the various up- and down-link fading patterns, and a given traffic load, a two-criteria optimization problem is faced. The decision variables are the service rates of the above-mentioned IP buffers for each SD pair, and the corresponding transmission parameters.

These allocation strategies are referred to as CLARA (Cross Layer Approach for Resource Allocation) in [17], where a few different criteria have been analyzed. In all cases, the performance evaluation of the system has been conducted on the basis of suitably defined indexes of the TCP connections’ goodput and the fairness of allocations. All CLARA techniques adopt a cross-layer optimization for the achievement of the best compromise between the maximization of TCP goodput and fairness for a combination of long-lived TCP Reno connections. The optimal allocations are numerically derived on the basis of the analytical model described in III.B, under different fade patterns, by using two methodologies. In the first one a balance is sought between the bandwidth allocations that maximize the overall goodput (the sum of

goodputs of all TCP connections over all SD pairs) and those that equalize all individual connections' goodputs, respectively; the corresponding optimal redundancy is also applied. Two different procedures to do this, named *range* and *tradeoff*, respectively, have been proposed. The allocations corresponding to the best choice of either procedure for each possible operating condition have also been considered, with the resulting strategy termed *merge*. The second methodology considers the maximization of the sum of the logarithms of the connections' goodputs, optimized with respect to the redundancy for each value of bandwidth allocation. This corresponds to seeking a Nash Bargaining Solution (NBS) over the connections, which has intrinsic fairness properties; the resulting allocation has been termed *Generalized Proportionally Fair (GPF)*. The different strategies have been compared in a static fading scenario first, and then in a dynamically varying one, with fading traces taken from real-life samples [17]. In this case, the allocation is applied adaptively, following the fading and traffic variations, similarly to the open-loop feedback strategies considered in the previous Section.

VI. A NUMERICAL EXAMPLE FOR TCP FLOWS

The numerical results we present in this Section only refer to the CLARA techniques mentioned in Section V. With respect to the results reported in [17], which were obtained with the same system parameters, we emphasize here statistical indexes of goodput and fairness, referring to the overall system. In order to apply the allocations in a dynamically varying environment, real fade traces are used. Table I reports the most significant system's parameters. In order to compute the link budget, we considered a portion of the Ka band (20/30 GHz) transponder of the Eutelsat satellite HotBird 6; data are derived from the file "Hot Bird 6 data sheet.fm", which is downloadable from [41]. Table II shows the configuration of the dynamic tests carried on; 10 different fading classes are considered, each one aggregating a certain number of connections.

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

System Parameter	Value
Stations' antenna diameter	1.2 m
Stations' maximum transmission 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 ($C / N_0 = 77.5$ dBHz) (10Mbit/s)	7.5 dB
Possible data coding rates	7/8 (clear sky), 3/4, 1/2
BER in clear sky after Viterbi decoder (7/8)	10^{-7}
mbu (min. bandwidth unit) in clear sky	5 kbit/s
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 mbu

TABLE II
CONFIGURATION OF THE TESTS

Class number	1	2	3	4	5	6	7	8	9	10
Number of Connections	2	3	2	4	2	3	1	4	3	6

For each configuration of C/N_0 values appearing in the fade traces for the various classes, the optimal allocations (bandwidth and redundancy) that correspond to the previously mentioned criteria are computed off-line, by means of the TEAM (TCP Elephant bandwidth Allocation Method) software, developed for this purpose⁽¹⁾. The results of the computations are then used in ns-2 simulation, by applying the optimal bandwidth/redundancy (i.e. bandwidth/packet loss) pairs that correspond to the time-varying fade traces.

Two performance indexes are used: the goodput factor and the fairness factor. The goodput factor represents the ratio of the overall instantaneous goodput under a generic

⁽¹⁾ The source code is available at the web address reported in [42].

allocation to the goodput-maximizing choice; the fairness factor measures the deviation of the allocation from the goodput-equalizing one. In particular, we defined:

$$\text{Goodput Factor: } \varphi_g = \frac{\sum_{i=1}^F n_c^{(i)} \hat{T}_g^{(i)}(B_i, r_i)}{\sum_{i=1}^F n_c^{(i)} \hat{T}_g^{(i)}(B_i^*, r_i^*)} \quad (3)$$

where F is the number of classes, $n_c^{(i)}$ is the number of connections in class i , (B_i, r_i) , $i = 1, \dots, F$, is any set of bandwidth-redundancy pairs, and (B_i^*, r_i^*) , $i = 1, \dots, F$, is the set that maximizes the global goodput.

$$\text{Fairness Factor: } \varphi_f = 1 - \frac{\sum_{j=1}^L |\hat{T}_g^{(j)} - \bar{T}_g|}{2\bar{T}_g(L-1)} \quad (4)$$

where $L = \sum_{i=1}^F n_c^{(i)}$ 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.

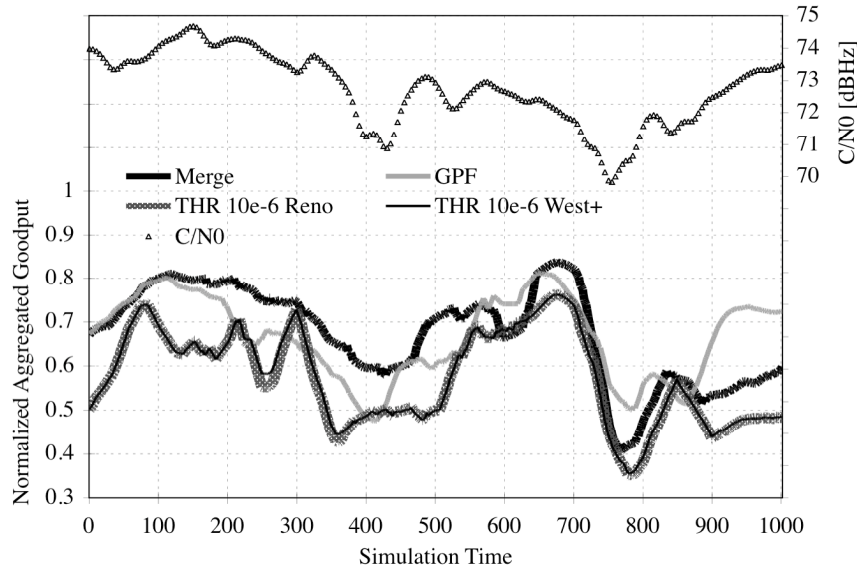


Fig. 1. Sample traces over a 1,000 s simulation interval of the overall goodput (averaged over a 10 sample moving window), normalized to the satellite uplink bandwidth, under different allocation strategies.

Figure 1 shows the behaviour of the normalized aggregated goodput over all classes, in a 1,000 s time window out of a total simulation time of 3,000 s. The “instantaneous” goodput is determined by the dynamic bandwidth and redundancy allocation, which

aims at countering the fading effects, and achieving a compromise between maximizing the total goodput and maintaining fairness among the connections. The *BER threshold* strategy is reported for comparison: it simply adjusts the redundancy to always keep the BER below a given limit, and assigns the bandwidth shares proportionally to the redundancy and to the number of connections of each class. Besides doing so with TCP Reno, which fits our model, we added for comparison also TCP Westwood+ [43] with BER threshold adaptation (labeled “THR 10e-6 West+” in Fig. 1). In the BER Threshold strategy, the BER threshold has been set to 10^{-6} , while in the Merge strategy the fairness factor threshold has been set to 0.85. Each goodput value is obtained by a moving average over a 10 s interval. The values of the C/N_0 curve in the figure are averaged over all classes.

In Fig. 2, we reported the relative frequency of various goodput and fairness factors, calculated in correspondence with each value of C/N_0 used in the simulation.

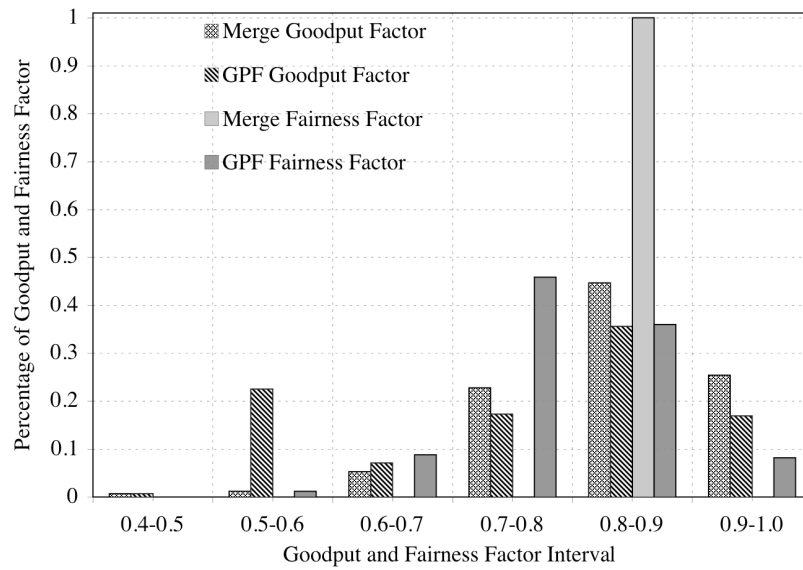


Fig. 2. Relative frequency of goodput and fairness factors.

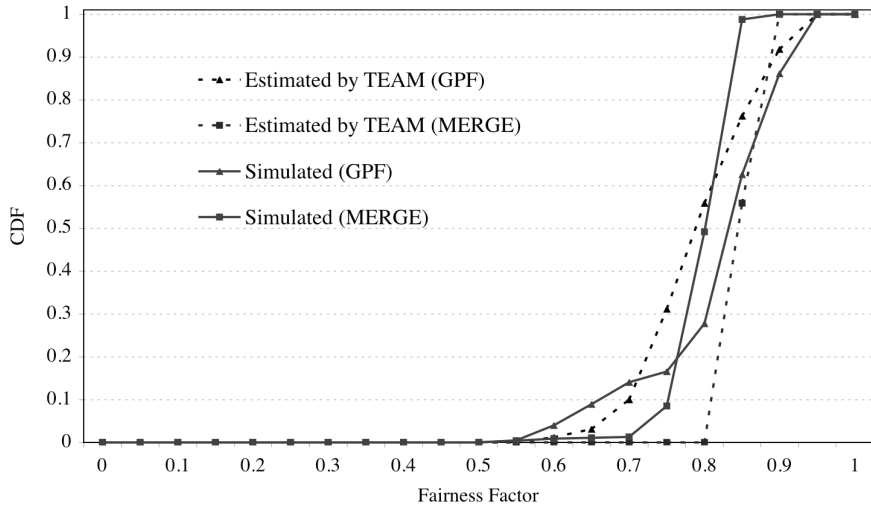


Fig. 3. Distribution of the fairness factors of Merge and GPF allocations, as computed by the TEAM software and by results of simulations, respectively.

In order to complete the comparison, Fig. 3 shows the Cumulative Distribution Function (CDF) of the fairness factor, stemming both from the TEAM calculation and from the output of simulations. The results show that in both Merge and GPF strategies the application of allocations to the real TCP allows to attain values that are not too far from those predicted by the model. The discrepancies are mainly due to two reasons, the first of which is the most relevant one: i) the values computed by the TEAM software refer to steady-state TCP behaviour, as would result if the fading configuration lasted forever, whereas the simulation applies allocations to a continuously evolving system; ii) calculations rely on a good but nevertheless approximated model, rather than on the detailed TCP representation used in simulation.

TABLE III
TOTAL GOODPUT AVERAGED OVER ALL CLASSES

Strategy		Goodput [Mbit/s]	Gain %
Merge		5.921	33.39
Range		5.858	31.97
Tradeoff		5.437	22.48
GPF		5.412	21.91
	BERthr10e-5	4.881	9.96
Westwood+	BERthr10e-6	4.626	4.21
	BERthr10e-7	4.617	4.00
	BERthr10e-5	4.428	-0.25
Reno	BERthr10e-6	4.619	4.05
	BERthr10e-7	4.439	0.00

A comparison among the goodput values achieved by various optimized cross-layer strategies of CLARA (*Merge, Range, Tradeoff, GPF*), averaged over all classes and all simulation time (3,000 s) has been performed in [17], also including TCP Westwood+ with BER threshold adaptation. For the reader's convenience, the results of the comparison are reported in Table III.

The gain in goodput, with respect to the last row, achieved when using TCP Reno is superior to that obtained when using TCP Westwood+, for all values of the BER threshold. This happens also in cases where Westwood+, which is more robust than Reno with respect to the BER, takes advantage of the higher achievable IBR, in the presence of less stringent BER constraints (10^{-5}). Thus, the cross-layer optimization (tailored on Reno, because of the model used) guarantees a performance that is superior to that of a modified TCP, even working under a BER threshold that well suits its characteristics. It is noted in [17] that, by adopting suitable models for the goodput, the application of a similar methodology to TCP Westwood+ could be investigated, in order to further enhance the performance. However, one of the main reasons to keep working with Reno) is that no prior modification is necessary on the end devices.

It is worth noticing that we chose the Reno version of TCP for our studies. Other versions, such as TCP NewReno and TCP SACK, are recognized to perform better over GEO satellite links in that they allow recovering multiple losses per RTT. However, in the maximization of the goodput, the field of interest is relative to high TCP efficiency and in these conditions we found (by using ns2 simulations) that all the three versions have practically the same behaviour.

VII. CONCLUSIONS

We have examined some RRM problems in multi-service satellite networks, by adopting a general cross-layer approach. The quantity of interest, to be dynamically shared among user stations and applications, is the up-link satellite bandwidth. In all cases considered there is a direct or indirect interaction between adjacent or even non-adjacent layers in the protocol architecture. In most cases, the effect of the physical layer, where fade countermeasures (in terms of bit and coding rate adaptation) are

applied to contrast fading attenuation, is taken into account as a reduction in the net available bandwidth. A few optimization problems have been addressed, even with different specific techniques, regarding MAC-physical, network-MAC and transport-physical interactions. Mixed RB and BE traffic classes and elastic TCP connections are the two main traffic scenarios considered, to be transmitted over a satellite channel that suffers from different fading conditions. We have presented a numerical example to highlight some characteristics of the dynamic cross-layer allocation in the case of long-lived TCP connections. One main purpose of the paper is to highlight the role of the redundancy coefficients in representing the effect of a specific fade countermeasure (code and rate adaptation) on CL-RRM. By explicitly taking this effect into account in the optimization problems that can be formulated for the bandwidth allocation (and related controls, like CAC), significant gains can be obtained, with respect to solutions that do not exploit this information.

REFERENCES

- [1] Hills A, Friday B. Radio resource management in wireless LANs. *IEEE Communications Magazine* 2004; **42**(12):S9-14.
- [2] Karetos GT, Kyriazakos SA, Groustiotis E, Di Giandomenico F, Mura I. A hierarchical radio resource management framework for integrating WLANs in cellular networking environments. *IEEE Wireless Communications Magazine* 2005; **12**(6):11-17.
- [3] Feng N, Mau S-C, Mandayam NB. Pricing and power control for joint network-centric and user-centric radio resource management. *IEEE Transactions on Communications* 2004; **52**(9):1547-1557.
- [4] Barsocchi P, Celandroni N, Ferro E, Gotta A, Davoli F, Giambene G, González Castaño FJ, Moreno JI, Todorova P. Radio resource management across multiple protocol layers in satellite networks: a tutorial overview. *International Journal of Satellite Communications and Networking* 2005; **23**(5):265-305.
- [5] Karaliopoulos M, Narenthiran K, Evans B, Henrio P, Mazzella M, de Win W, Dieudonne M, Philippopoulos P, Axiotis DI, Andrikopoulos I, Mertzanis I, Corazza GE, Vanelli-Coralli A, Dimitriou N, Polydoros A. Satellite radio interface and radio resource management strategy for the delivery of multicast/broadcast services via an integrated satellite-terrestrial system. *IEEE Communications Magazine* 2004; **42**(9):108-117.
- [6] Roberts JW. A survey on statistical bandwidth sharing. *Computer Networks* 2004; **45**(3):319-332.
- [7] Roberts JW. Internet traffic, QoS and pricing. *Proceedings of the IEEE* 2004; **92**(9):1389-1399.

- [8] Kawadia V, Kumar PR. A cautionary perspective on cross layer design. *IEEE Wireless Communications Magazine* 2005; **12**(1):3-11.
- [9] Alouf S, Altman E, Galtier J, Lalande J-F, Touati C. Quasi-optimal resource allocation in multi-spot MFTDMA satellite networks. In Cheng M, Li Y, Du D Z, Eds. *Combinatorial Optimization in Communication Networks*, Kluwer Academic Publishers, 2005; 1-41.
- [10] Park J-M, Savagaonkar U, Chong EKP, Siegel HJ, Jones SD. Allocation of QoS connections in MF-TDMA satellite systems: a two-phase approach. *IEEE Transactions on Vehicular Technology* 2005; **54**(1):177-190.
- [11] Alagöz F, Vojcic BR, Walters D, AlRustamani A, Pickholtz RL. Fixed versus adaptive admission control in direct broadcast satellite networks with return channel systems. *IEEE Journal on Selected Areas in Communications* 2004; **22**(2):238-249.
- [12] Açar G, Rosenberg C. Simulation analyses of weighted fair bandwidth-on-demand (WFBoD) process for broadband multimedia geostationary satellite systems. *International Journal of Satellite Communications and Networking* 2005; **23**(4):229-245.
- [13] Jiang Z, Leung VCM. A predictive demand assignment multiple access protocol for Internet access over broadband satellite networks. *International Journal of Satellite Communications and Networking* 2003; **21**(4-5):451-467.
- [14] Chiti F, Fantacci R, Tarchi D, Kota S, Pecorella T. QoS provisioning in GEO satellite with onboard processing using predictor algorithms. *IEEE Wireless Communications Magazine* 2005; **12**(5):21-27.
- [15] Kota S, Marchese M, Quality of Service for satellite IP networks: a survey. *International Journal of Satellite Communications and Networking* 2003; **21**(4-5):303-348.
- [16] Celandroni N, Davoli F, Ferro E, Gotta A. An overview of some techniques for cross-layer bandwidth management in multi-service satellite IP networks. In *Proceedings of the Workshop on "Advances in Satellite Communications: New Services and Systems"*, IEEE Globecom '05, St. Louis, MO, Nov.-Dec. 2005; WO4:4.1-WO4:4.6.
- [17] Celandroni N, Davoli F, Ferro E, Gotta A. Long-lived TCP connections via satellite: cross-layer bandwidth allocation, pricing and adaptive control. *IEEE/ACM Transactions on Networking* 2006; to appear.
- [18] Celandroni N, Davoli F, Ferro E, Gotta A. Adaptive cross-layer bandwidth allocation in a rain-faded satellite environment. *International Journal of Communication Systems* 2006; to appear [Published Online: 14 Nov 2005; DOI: 10.1002/dac.751 - <http://www3.interscience.wiley.com/cgi-bin/jhome/5996>].
- [19] Celandroni N, Davoli F, Ferro E. Static and dynamic resource allocation in a multiservice satellite network with fading. *International Journal of Satellite Communications and Networking* 2003; **21**(4-5):469-487.
- [20] Iuoras N, Le-Ngoc T. Dynamic capacity allocation for quality-of-service support in IP-based satellite networks. *IEEE Wireless Communications Magazine* 2005; **12**(5):14-20.

- [21] Le-Ngoc T, Leung V, Takats P, Garland P. Interactive multimedia satellite access communications. *IEEE Communications Magazine* 2003; **41**(7):78-85.
- [22] Durrezi A, Jagannathan PK, Jain R. Proportional bandwidth allocation in satellite networks. In *Proceedings of the 2003 IEEE Aerospace Conference*, Big Sky, MT, March 2003; 1253-1264.
- [23] Neely MJ, Modiano E, Rohrs CE. Power allocation and routing in multibeam satellites with time-varying channels. *IEEE/ACM Transactions on Networking* 2003; **11**(1):138-152.
- [24] Fu AC, Modiano E, Tsitsiklis JN. Optimal energy allocation and admission control for communications satellites. *IEEE/ACM Transactions on Networking* 2003; **11**(3):488-500.
- [25] Iera A, Molinaro A, Marano S. Traffic management techniques to face the effects of intrinsic delays in geostationary satellite networks. *IEEE Transactions on Wireless Communications* 2002; **1**(1):145-155.
- [26] Baglietto M, Davoli F, Marchese M, Mongelli M. Neural approximation of open loop–feedback rate control in satellite networks. *IEEE Transactions on Neural Networks* 2005; **16**(5):1195-1211.
- [27] Castanet L, Bolea-Alamanac A, Bousquet M. Interference and fade mitigation techniques for Ka and Q/V band satellite communication systems. In *Proceedings of the International 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>].
- [28] Digital Video Broadcasting, frame structure, channel coding and modulation for 11-12 GHz satellite services. *ETSI EN 300 421 V1.1.2* 1997.
- [29] Digital Video Broadcasting (DVB). Interaction channel for satellite distribution systems. *ETSI EN 301 790* 2003.
- [30] Ross KW. *Multiservice Loss Models for Broadband Telecommunication Networks*. Springer-Verlag: London, 1995.
- [31] Alagöz F. Approximations on the aggregated MPEG traffic and their impact on admission control. *Turkish Journal of Electrical Engineering and Computer Sciences* 2002; **10**(1):73–84.
- [32] Tsybakov B, Georganas ND. Self-similar traffic and upper bounds to buffer-overflow probability in an ATM queue. *Performance Evaluation* 1998; **32**(1):57-80.
- [33] Bertsekas DP. *Dynamic Programming and Optimal Control*. Athena Scientific: Belmont, MA, 1995.
- [34] Ebrahimi-Taghizadeh S, Helmy A, Gupta S. TCP vs. TCP: A systematic study of adverse impact of short-lived TCP flows on long-lived TCP flows. In *Proceedings of IEEE INFOCOM '05*, Miami, FL, March 2005; 926-937.
- [35] Padhye J, Firoiu V, Towsley DF, Kurose JF. Modeling TCP Reno performance: a simple model and its empirical validation. *IEEE/ACM Trans. Networking* 2000; **8**(2):133-145.
- [36] Celandroni N, Potorti F. Maximising single connection TCP goodput by trading bandwidth for BER. *International Journal of Communication Systems* 2003; **16**(1):63-79.

- [37] Marchese M, Mongelli M. On-line bandwidth control for quality of service mapping over satellite independent service access points. *Computer Networks* 2006; in press. Published on line at <http://www.sciencedirect.com/science/journal/13891286>.
- [38] ETSI, Satellite Earth Stations and Systems (SES). Broadband Satellite Multimedia, Services and Architectures. *ETSI Technical Report*, TR 101 984 V1.1.1, Nov. 2002.
- [39] ETSI, Satellite Earth Stations and Systems (SES). Broadband Satellite Multimedia, IP over Satellite, *ETSI Technical Report*, TR 101 985 V1.1.2, Nov. 2002.
- [40] Cassandras CG, Sun G, Panayiotou CG, Wardi Y. Perturbation analysis and control of two-class stochastic fluid models for communication networks. *IEEE Transactions on Automatic Control* 2003; **48**(5):23-32.
- [41] <http://www.eutelsat.com/satellites/13ehb6.html>
- [42] <http://www.isti.cnr.it/ResearchUnits/Labs/wn-lab/software-tools.html>
- [43] Dell'Aera A, Grieco LA, Mascolo S. Linux 2.4 implementation of Westwood+ TCP with rate-halving: a performance evaluation over the Internet. In *Proceedings of the IEEE International Conference on Communications 2004 (ICC 2004)*, Paris. France, June 2004; 2092-2096.