

TCP performance measured over heterogeneous wireless networks with high delay-bandwidth product*

Antonio Annes[~], Nedo Celandroni[#], Franco Davoli[^], Erina Ferro[#], Alberto Gotta[#],

[~]CISCO, Bruxelles (Belgium), aannese@cisco.com

[#]CNR-ISTI, Via Moruzzi 1, Pisa (Italy) {nedo.celandroni, erina.ferro, alberto.gotta}@isti.cnr.it

[^]CNIT, University of Genoa Research Unit, Via dell'Opera Pia 13, Genoa (Italy),
franco.davoli@cnit.it

Abstract—The performance of three variants of TCP (Westwood+, NewReno, and NewReno with SACK option), available in the Linux Kernel, have been measured in a heterogeneous satellite-terrestrial wireless integrated network, characterized by high delay-bandwidth product and random losses. The testbed is made up of a set of portable computers, interconnected by WLANs, and a remote host connected via a geostationary satellite link. The target of this work is to measure the real performance of the TCP protocol, in order to verify the accuracy of some models proposed in the literature. In fact, most of these models have generally been validated only by means of analytical or simulation approaches, but without the use of real measurements. We utilized the SkyplexData commercial satellite platform, which adopts routers characterized by very large buffers. The *goodput* of both a single and two simultaneous TCP flows has been measured. Moreover, timeouts' occurrences, which produce the most relevant effects in such an environment, are thoroughly investigated.

Index Terms—Satellite, wireless network, TCP, goodput, measures

I. INTRODUCTION

Wireless networks are an easy, fast, and flexible way to provide the first communication means for scenarios where a wired infrastructure is lacking or is not economically feasible. In these situations, Mobile Ad hoc NETWORKS (MANETs) [1], interconnected through a backbone satellite link, constitute a possible solution for an entirely wireless platform, in which hosts/clients communicate with one or more servers.

The Transmission Control Protocol (TCP) [2] is currently the most widely used transport protocol in packet networks [3], such as the Internet, and it is largely responsible for end-to-end congestion control. The congestion control is obtained by

throttling the rate at which a source injects packets into the network; this is carried out by reducing the transmission rate at any perceived congestion event, and by gradually increasing this rate in between two consecutive congestion events. It is well known that the loss recovery algorithm is not aware of the causes of a packet loss; thus, TCP cannot distinguish between packet corruptions and losses due to congestion. The result is that TCP reduces its transmission rate in both cases; this is realized in classical TCP versions by halving the congestion window. Due to the nature of its congestion control mechanism, TCP Reno suffers from performance problems when multiple packets are lost in a window of data [4]. In a MANET-satellite scenario, high latency, bandwidth asymmetry, and transmission errors on wireless channels remarkably affect the TCP performance. Regarding the terrestrial wireless hop, channel losses due to (time-varying) multipath fading may be dominant (or comparable to) congestion losses.

In recent years, enhancements of TCP standard mechanisms significantly improved the performance of long-lived connections with multiple random losses in the same window of data. NewReno addresses this problem by introducing *partial acknowledgments* in the loss recovery phase [5], while




Fig. 1. The test-bed configuration

SACK [6] adopts a *selective acknowledgment* mechanism, combined with a selective repeat retransmission policy.

*Work funded by the European Commission in the framework of the "SatNEx" NoE (contract no. 507052), and partially by MIUR (Italy) in the framework of the IS-MANET and CRIMSON projects.

Westwood+ [7] dynamically estimates the available bandwidth to set the congestion window and the slow start threshold after a packet loss.

The literature on TCP modeling is rather vast [8-11] and has been recently targeted toward wireless communications [12-16]. Such models are sometimes derived under restrictive and not very realistic assumptions, which may be far from the real experience. In this paper, we aim at clarifying the real causes that may affect the functioning of TCP over terrestrial wireless networks integrated with a satellite backbone or, more generally, in networks with large delay-bandwidth product paths. Moreover, we compare the performance of the three above-mentioned TCP variants by means of measurements performed over a real satellite platform, which includes a

TABLE I
SKYPLEX PARAMETERS AND LINUX KERNEL VARIABLE SETTINGS

Parameter/variables	Value
Min measured RTT	557 [ms]
Mean measured RTT	701 [ms]
Gross Bandwidth	2 [Mb/s]
Net Bandwidth	1.2 [Mb/s]
Buffer size on Tx modem	1500 [bytes]
net.ipv4.tcp_window_scaling	1
net.ipv4.tcp_rmem (min, default, max)	8388608 8388608 8388608 [bytes]
net.ipv4.tcp_wmem (min, default, max)	8388608 8388608 8388608 [bytes]
net.ipv4.tcp_mem (min, default, max)	8388608 8388608 8388608 [bytes]
net.core.rmem_default	8388608 [bytes]
net.core.wmem_default	8388608 [bytes]
net.core.rmem_max	8388608 [bytes]
net.core.wmem_max	8388608 [bytes]

MANET; the test-bed utilized is available at the Wireless Laboratory of CNR-ISTI in Pisa (Italy).

II. THE TEST BED PLATFORM

The test-bed configuration used in our transmission experiments is depicted in Figure 1.

The MANET is constituted by nine notebooks (Celeron 1.133 GHz with 256 MB of RAM), each equipped with a Debian Linux operating system (kernel 2.6.11) with Web100 kernel patch [17]. PCMCIA network cards (Conceptronic C54RC IEEE802.11g for transmission and reception) realize the terrestrial wireless link; they always operate with the standard DCF (Distributed Coordination Function) [18] contention access mode at a rate of 11 Mbps. Each notebook can act as a traffic source. A Linux host acts as a gateway between the WLAN and the satellite network, and forwards packets towards the satellite. The satellite link is a Ka channel of the Eutelsat HotBirdTM6 spacecraft, which uses the commercial technology known as Skyplex Data [19], [20]. This is an IP-based satellite network derived from DVB-RCS [21], and it is used here in a standard point-to-point topology (for unicast traffic). We used two remote Linux hosts, located at the CNIT National Laboratory for Multimedia Communications in

Naples and at the CNIT Research Unit in Pisa, respectively, as traffic sinks at the other side of the satellite link. In unicast transmissions, WLAN cards use standard IEEE802.11 settings; in particular, the DCF Retry Limit value is set to 7 for every outgoing MAC frame, i.e., up to 7 transmission retries are performed in case of unacknowledged frames. Table I

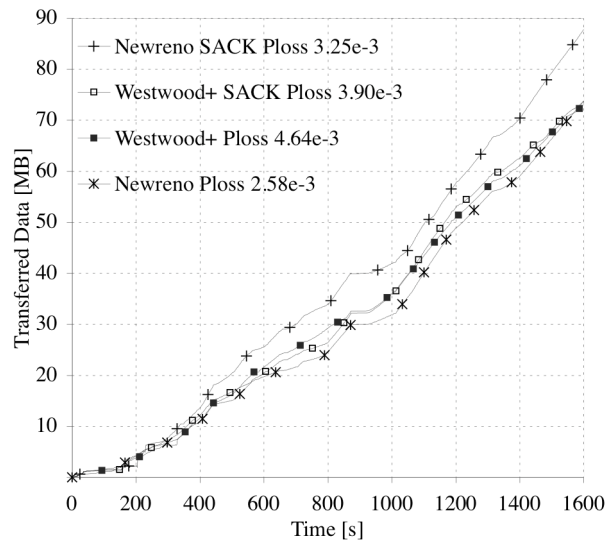


Fig. 2. Transferred data vs. Time with four different TCP variants. shows some of the Skyplex parameters and the Kernel configuration we used.

III. EXPERIMENTAL RESULTS

We focused the attention on TCP NewReno, and on the TCP Westwood+ variant, both with and without the SACK option, since all these implementations are available on recent Linux Kernels. We remind here that the SACK option requires the receiver be aware of the enabled option. The trials have been performed with two concurrent TCP connections, one per variant, for a duration of 1600s per trial. The satellite link was error-free, while the wireless hop experienced a packet loss, which is reported in Table II. The variables we monitored are the RTT (round trip time), the Congestion Window, and the Timeouts; furthermore, we measured the instantaneous *goodput* per connection and the Transferred Data versus time.

Some statistics of the collected data are shown in Table II, where only the most significant trials, among all those performed, are reported. The results of four trials are shown in the table; each trial compares two different TCP versions. In each trial two simultaneous flows (each one originated by the TCP versions used in the trial), share the same terrestrial wireless and satellite hops, as indicated in Fig.1.

The first two trials of Table II (rows 1-4) show the improvement in performance of NewReno and Westwood+, when the SACK option is active. Trials 3 and 4 (rows 5-8) compare NewReno with Westwood+ with and without the SACK option, respectively.

Figure 2 shows the amount of data transferred over time for all the four variants considered (corresponding to connections that belong to trials # 1 and 2). We have also measured the number of packets lost on the WLAN hop, by means of a

TABLE II
 MEAN AND MEDIAN OF: RTT, CONGESTION WINDOW, AND NORMALIZED GOODPUT; AVERAGE PACKET LOSS; NUMBER OF TIME OUTS (T.O.)

Trial Number	TCP variant	Average Packet Loss	Mean, Median			Number of timeouts
			RTT [s]	CWND [MSU=1500B]	Normalized Goodput	
1	Newreno	2.58e-3	1.56, 1.23	59.6, 50.2	0.32, 0.38	16
	Newreno SACK	3.25e-3	2.03, 1.37	119.0, 63.0	0.39, 0.42	11
2	Westwood+	4.64e-3	1.83, 1.24	101.0, 54.0	0.31, 0.33	21
	Westwood+ SACK	3.90e-3	1.72, 1.09	78.0, 45.0	0.32, 0.32	12
3	Newreno	1.43e-3	1.90, 1.30	58.9, 55.0	0.27, 0.28	21
	Westwood+	1.10e-3	1.65, 1.34	58.3, 56.0	0.30, 0.28	11
4	Newreno SACK	3.11e-4	3.99, 3.99	196, 170	0.46, 0.45	1
	Westwood+ SACK	4.32e-4	4.05, 4.10	213, 224	0.50, 0.52	2

packet *sniffer*. This parameter is very important, in order to draw some considerations on this comparison. NewReno SACK transfers the largest amount of data even if the experienced packet loss is the second highest value.

The trial with Westwood+ reported in Fig. 2 presents the highest value of packet loss; nevertheless, the volume of its transferred data is comparable with the other ones, which

experience fewer loss events.

Figure 3 shows the instantaneous *goodput* measured in the four trials. The sender TCP agent computes this quantity at each ACK reception, by dividing the segment size by the inter-ACK time. This explains the big *jitter* exhibited in the charts. Let us consider trial number 4, with two parallel flows

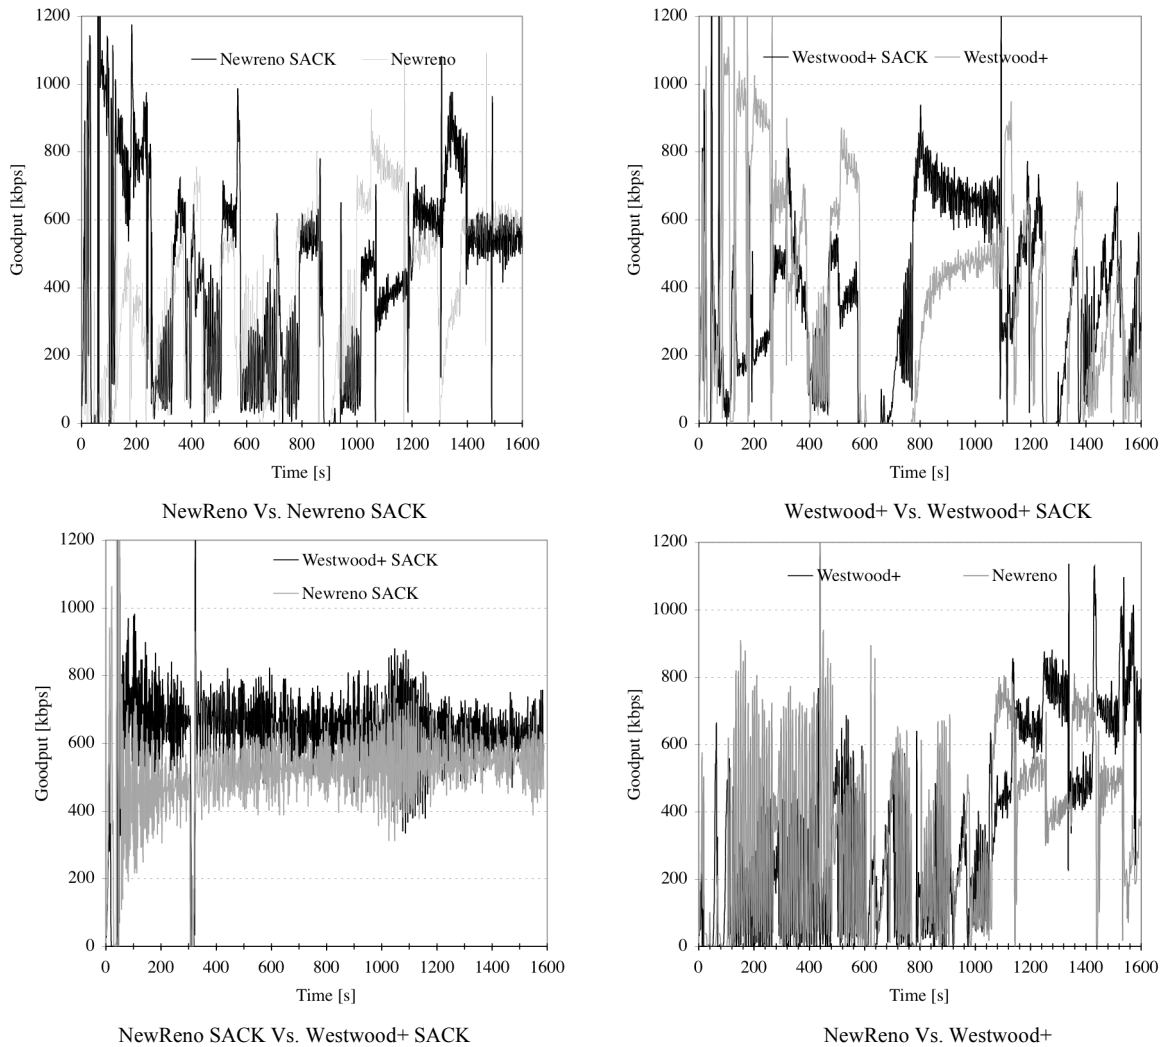


Fig 3 – Instantaneous Goodput of two TCP connections, which share the same wireless link in four different trials.

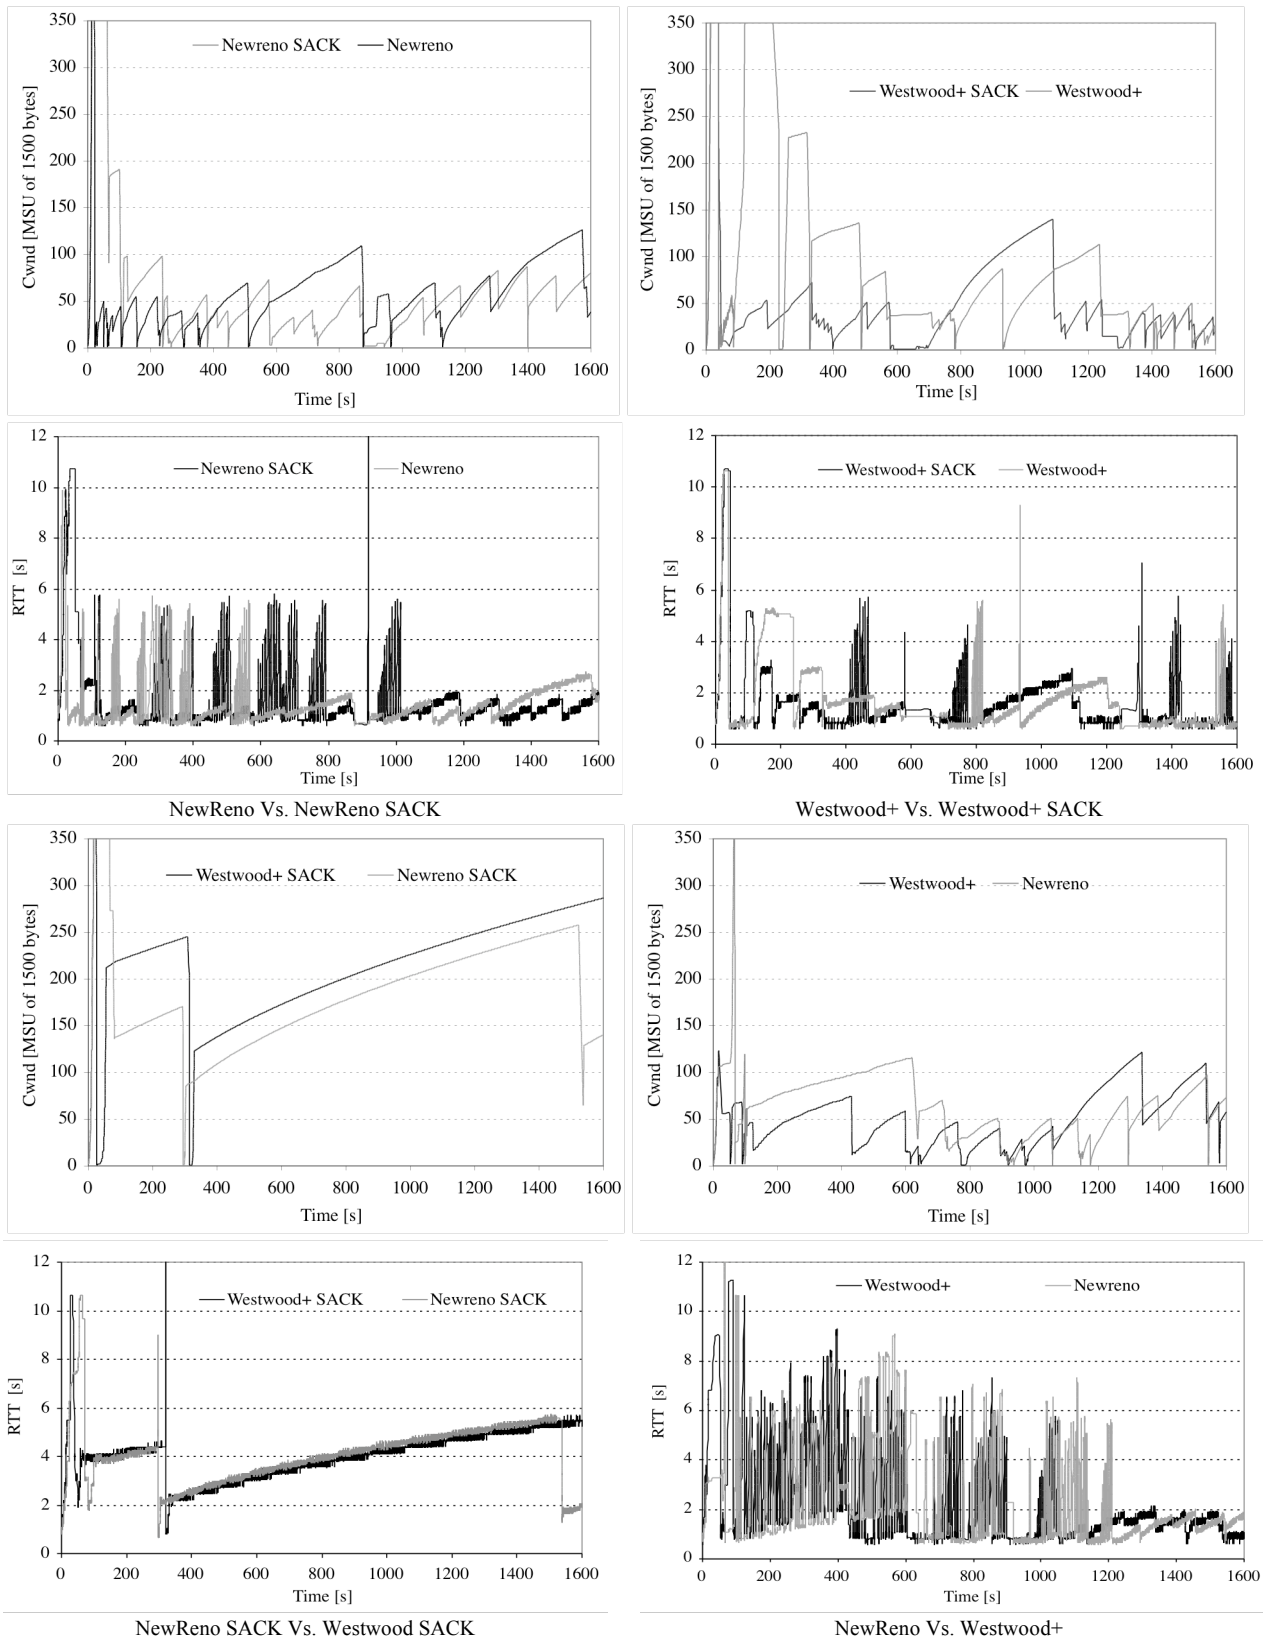


Fig 4 – Congestion window and round trip time of two TCP connections, which share the same wireless link in four different trials.

of NewReno and Westwood+ both with SACK option. In such a case with very few loss events (see Table II), Westwood+ exhibits unfairness with respect to NewReno in the bandwidth

sharing, resulting in about 15% more than a fair sharing.

We have also measured the round trip time and the congestion window values, getting these values from the kernel variables, during the aforementioned trials. The behaviour of these variables versus time is shown in Fig. 4. Looking at the round trip time's traces, we observe that a large jitter, even longer than 5 seconds, sometimes occurs. In correspondence of these effects, the congestion windows are reduced to a single packet, when a timeout occurs. These effects are due to a temporary outage of the terrestrial wireless link. In conditions of poor signal-to-noise ratio, we observed that the wireless link went down even for several seconds. Thus, the RTT has been momentarily increasing for the duration of the outage.

With respect to this fact, we remark the difficulty in finding a proper positioning of the terminals that would produce an almost stable situation in packet loss.

The aim of this study did not include the explicit consideration of layer II recovery algorithms, which recover errors on the 802.11 link. Thus, we let the wireless cards choose their transmission rate, as happens customarily in every-day usage. On the other hand, we try to study the effects of a significant number of losses on the TCP behaviour over a Fat Lossy Channel. With very little variations of the signal level, the error recovery technique, applied at layer II, may cause a threshold effect, where either the largest part of transmitted packets are recovered (as in trial 4) or the link goes down for an outage period.

IV. CONCLUSIONS

The behaviour of different versions and implementations of the TCP protocol has been studied over a real test bed, in a situation characterized by the presence of a satellite channel (introducing long delay) and a WLAN segment (introducing channel errors). In particular, the specific satellite platform used (the CNIT satellite network, provided by Eutelsat) adopts relatively large IP buffers (in the order of 1.5 MB) in the access routers, thereby increasing the RTT variability. All versions have been tested in pairs (one connection per type), over the same path, in order to also verify the fairness of one with respect to the other, and to operate under the same fading pattern in both cases. With respect to the fairness, Westwood+ with SACK exhibits a more aggressive behaviour towards NewReno with SACK, whereas all other combinations are relatively well balanced. Westwood+ appears to be more robust with respect to packet losses, and NewReno with SACK achieves the highest normalized *goodput* under relatively more severe fading conditions (see trials # 1 and 2 in Table II). It should be noted that the SACK option must be agreed upon at both transmitter and receiver, whereas Westwood+ needs a modification in the transmitter only.

ACKNOWLEDGMENT

We are grateful to Mr. Antonio Blasco Bonito and Mr. Alessio Gemma from CNR-ISTI, who helped us in developing and testing the network topology used for the measurement

campaign. We are likewise very grateful to Pisa and Naples CNIT Research Units, for making available the remote hosts on the other side of the satellite link.

REFERENCES

- [1] S. Corson, J. Macker, "Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations", RFC 2501, Jan. 1999.
- [2] M. Allman, V. Paxson, W. S. Stevens. TCP Congestion Control. RFC 2581, April 1999.
- [3] J. Bolot, "End-to-end packet delay and loss behavior in the Internet," *Proc. ACM SIGCOMM'93*, San Francisco, CA, Aug. 1993, pp. 289–298.
- [4] K. Fall, S. Floyd, "Comparisons of Tahoe, Reno, and Sack TCP," manuscript, <ftp://ftp.ee.lbl.gov>, Mar. 1996.
- [5] S. Floyd, T. Henderson, New Reno Modification to TCP's Fast Recovery, RFC 2582, April 1999.
- [6] M. Mathis, J. Mandavi, S. Floyd, A. Romanow, "TCP Selective Acknowledgement Options", RFC 2018, Oct. 1996.
- [7] S. Mascolo, L. A. Grieco, R. Ferorelli, P. Camarda, G. Piscitelli, "Performance evaluation of Westwood+ TCP congestion control", *Performance Evaluation*, vol. 55, no. 1-2, pp. 93-111, Jan. 2004.
- [8] A. Kumar, "Comparative performance analysis of versions of TCP in a local network with a lossy link," *IEEE/ACM Trans. Networking*, vol. 6, no. 4, pp. 485–498, 1998.
- [9] J. Padhye, V. Firoiu, D. Towsley, J. 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.
- [10] M. Mathis, J. Semke, J. Madhavi, T. Ott, "The macroscopic behavior of the TCP congestion avoidance algorithm," *Comput. Commun. Rev.*, vol. 27, no. 3, pp. 67–82, 1997.
- [11] J. Padhye, V. Firoiu, D. Towsley, "A stochastic model of TCP Reno congestion avoidance and control." Univ. Massachusetts, Cambridge, MA, *Tech. Rep. CMPSCI-99-02*, 1999.
- [12] E. Altman, K. Avrachenkov, C. Barakat, "A stochastic model of TCP/IP with stationary random losses," *Proc. ACM SIGCOMM'2000*, Stockholm, Sweden, Aug. 2000.
- [13] T. 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.
- [14] M. Zorzi, A. Chockalingam, R. R. Rao, "Throughput analysis of TCP on channels with memory," *IEEE J. Select. Areas Commun.*, vol. 18, no. 7, pp. 1289–1300, July 2000.
- [15] F. Anjum, L. Tassiulas, "On the behavior of different TCP algorithms over a wireless channel with correlated packet losses," *Proc. ACM SIGMETRICS'99*, Atlanta, GA, May 1999, pp. 155–165.
- [16] A. A. Abouzeid, S. Roy, M. Azizoglu, "Comprehensive performance analysis of a TCP session over a wireless fading link with queueing", *IEEE Trans. Wireless Commun.*, vol. 2, no. 2, pp. 344–356, March 2003.
- [17] www.web100.org
- [18] IEEE 802.11. *Information Technology Telecommunications and Information Exchange between Systems Local and Metropolitan Area Networks Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. ANSI/IEEE Std. 802.11, ISO/IEC 8802-11, first edition, 2001.
- [19] C. Elia, E. Colzi. "Skyplex: distributed up-link for digital television via satellite", *IEEE Trans. Broadcasting*, vol. 42, no. 4, pp. 305-313, Dec. 1996.
- [20] http://www.eutelsat.com/fr/satellites/pdf/dealers/annex_d_skyplex.pdf
- [21] EN 301 192 v1.2.1, *Digital Video Broadcasting (DVB): Specification for data broadcasting*, June 1999