PACKET LOSS IN TCP HYBRID WIRELESS NETWORKS

Paolo Barsocchi, Gabriele Oligeri, and Francesco Potortì

***Abstract*—Hybrid networks where packets travel through both a geostationary satellite and a terrestrial local wireless link are a particularly challenging environment with respect to TCP performance. The reason is that such networks exhibit both a high latency and a high packet loss, a combination that stresses TCP's congestion avoidance algorithm. Accurate modelling of the packet loss process in such networks is not an easy task, so very simplified models are frequently used. We start from channel state measurements done on a fixed-speed IEEE 802.11 terrestrial channel and we use the ns-2 simulator to evaluate the performance of TCP on it. We then evaluate the accuracy of some models that are commonly used to represent the channel state, by comparing these models with actual traces.**

I. Introduction

THE presence of both a geostationary satellite link and a terrestrial local wireless link on the same path of a given network connection is becoming increasingly common, thanks to the popularity of the IEEE 802.11 protocol [1]. In the following, we will refer to this ubiquitous wireless terrestrial protocol as Wi-Fi, the popular trademark of the association founded to set interoperability standards for it.

The most common situation where a hybrid network comes into play is having a Wi-Fi link at the network edge and the satellite link somewhere in the network core. Example of scenarios where this can happen are ships or airplanes where Internet connection on board is provided through a Wi-Fi access point and a satellite link with a geostationary satellite; a small office located in remote or isolated area without cabled Internet access; a rescue team using a mobile ad hoc Wi-Fi network connected to the Internet or to a command centre through a mobile gateway using a satellite link [2]-[3].

The serialisation of terrestrial and satellite wireless links is problematic from the point of view of a number of applications, be they based on video streaming, interactive audio or TCP. The reason is the combination of high latency, caused by the geostationary satellite link, and frequent, correlated packet losses caused by the local wireless terrestrial link. In fact, GEO satellites are placed in equatorial orbit at 36,000 km altitude, which takes the radio signal about 250 ms to travel up and down. Satellite systems exhibit low packet loss most of the time, with typical project constraints of 10^{-8} bit error rate 99% of the time, which translates into a packet error rate of 10^{-4}, except for a few days a year. Wi-Fi links, on the other hand, have quite different characteristics. While the delay introduced by the MAC level is in the order of the milliseconds, and is consequently too small to affect most applications, its packet loss characteristics are generally far from negligible. In fact, multipath fading, interference and collisions affect most environments, causing correlated packet losses: this means that often more than one packet at a time is lost for a single fading event [4].

In this paper, we concentrate on frame error models targeted to TCP-based applications, for which the combination of high, correlated packet loss and high latency may cause very bad performance. Since TCP interprets packet loss as a sign of congestion, thus slowing its pace to avoid worsening the network conditions, frame errors due to corruption causes decreased throughput, even in the absence of network congestion. Various techniques exist that tackle this problem [5], but no definitive answers yet. This is consequently a widely studied research topic, often by means of simulations, for a variety of reasons, related partly to the usual simulation advantages, such as repeatability, ease in changing the problem's parameters and statistical significance, and partly to the peculiarities of the environment, namely the fact that satellite links are expensive and not widely available, especially for high bit rates.

A complete packet error model requires three levels of operation: a *frame error* model at the raw channel level, an implementation of the Wi-Fi *ARQ* (automatic retransmission) mechanism, and an implementation of the vendor-dependent *speed switching* mechanism. The first level gives a synthetic statistical representation of the frame error process on the raw channel due to one of two effects: bad preamble acquisition and, for correctly acquired frames, bad CRC due to corrupted bits in the MAC protocol data unit. The second level implements the ARQ mechanism defined in the IEEE 802.11 standard, by which a transmitter retries sending a frame up to a configurable number of times – typically 7 – when it does not receive an acknowledgement for the frame sent. The third level implements an adaptive modulation and coding scheme, by which the sender may choose to switch to a different transmission speed when the perceived channel conditions change. This study tackles the first two levels out of the three that are required for a complete packet error

Manuscript abstract received January 31, 2006. This work was supported by the CNR/MIUR under Legge 449/97 (project IS-Manet) and by the European Commission under the European Satellite Communications NoE (SatNEx, IST-507052) within the 6th Research Framework Programme.

All authors are with the ISTI - CNR (Institute of the National Research Council), via Moruzzi 1, 56124 Pisa, Italy; e-mail: {Paolo.Barsocchi, Gabriele.Oligeri, Potorti}@isti.cnr.it.

model. It aims at choosing an adequate statistical model for the raw channel frame error by studying how TCP behaves on a fixed-speed channel with ARQ.

II. LINK CHARACTERISTICS

In the proposed hybrid scenario, two links with very different characteristics are traversed by the TCP flow. Each link, from the point of view of TCP, is defined by four characteristics:

- *capacity*, which is the available rate seen at the IP level;
- *latency*, that is the time a packet takes to traverse the link;
- *buffer* size, which is the size of the buffer that a router or bridge uses to store packets waiting to be sent on the link;
- *segment error*, which is a description of the process causing segment drops due to link errors.

A. Satellite link

As far as the *capacity* is concerned, we considered a 2 Mb/s link, which is a common commercial figure.

As far as the delay is concerned, the geostationary channel introduces a signal delay of around 250 ms, the exact number depending on the satellite's position with respect to the earth stations. To this, one should add the processing delays at the earth stations due to segmentation, interleaving, coding, and the reverse operations at the receiver's side. When TDMA schemes are used, there is delay due to the frame structure, and possibly delay due to dynamic allocation. Moreover, there are processing delays at the routers, so any *latency* not smaller than 300 ms is a reasonable assumption.

The routers at the earth stations have buffers on the incoming terrestrial interfaces, where traffic is stored prior to sending on the satellite link. This is necessary for TDMA systems, where packets must wait for the first opportunity to be transmitted, is used by dynamic allocation methods where the allocation can be based on traffic backlog at the earth stations, and is generally useful for absorbing traffic rate fluctuations. From the TCP point of view, moreover, the buffer has the important role of allowing complete utilisation of the link capacity. In fact, it can be shown that the channel can be fully used only if no errors are present and the buffer size is at least equal to $C*RTT$, where C is the channel capacity and RTT is twice the link latency. We set the *buffer* in our experiments to exactly the said value, that is, 150 kB.

Geostationary satellite links are commonly depicted as AWGN channels, that is, affected by additive white Gaussian noise. In fact, given the extremely small antenna aperture, the high pointing precision and the absence of mobility, the only relevant noise at the receiver is the thermal noise, which gives rise to uncorrelated bit corruption and, in turn, to uncorrelated frame errors with constant probability; for this reason, in clear sky, we assume a Bernoullian channel, sometimes called a Poisson channel.

In this paper, we neglect the issues related to signal fading due to variable atmospheric conditions, which would require a different treatment, and assume a satellite channel with a *FER* (Frame Error Rate) of 10^{-5} which, for our purpose, is the same as errorless, because it is about the same FER which is caused by congestion [6]. This figure is consistent with measurements done on a recent Skyplex-based satellite system [7].

B. Wi-Fi link

Wi-Fi links comply with the IEEE 802.11 [1] standard, which has provisions for *ACM* (Automatic Coding and Modulation). ACM, which we call *speed switching*, is widely implemented in commercial Wi-Fi devices, but the algorithm for choosing which speed to adopt is not defined by the standard and still the subject of research. In this paper we analyse the behaviour of TCP with different frame error traces at a fixed Wi-Fi nominal rate of 11 Mb/s. Our experiments show that real network cards behave as the standard dictates with very little error [4]. Computations based on the standard give a *capacity* of 5.3 Mb/s for a unidirectional flow of 1000-data byte packets, long preambles, no fragmentation and RTS/CTS disabled [8].

Our experiments show a latency produced by the driver and hardware around 1 ms, plus delays due to retransmissions greater than 50 ms in less than 0.1% of packets. From the TCP point of view, these delays are not significant, so we considered the *latency* of the WiFi link as null.

The link *buffer* has no influence on the TCP dynamics, as long as it can keep few back-to-back segments, because the bottleneck is on the satellite link. Any buffer size greater than four packets' worth would yield the same results.

As far as the *error* model is concerned, this is the topic of Section 4. Measurement traces and synthetic traces were used, and an ad hoc error model was developed for ns-2 called *Ttem* (time trace error model). Using this error model, we were able to emulate a Wi-Fi channel affected by noise, where the state of the channel (good or bad) is described by a trace file containing a 0 for a good interval where no error happens and a 1 for a bad interval where a frame is corrupted; we made measurements at 5 ms long intervals. The Ttem error model keeps into account the IEEE 802.11 ARQ algorithm, but no speed switching. Ttem, as a patch to ns-2.29, is available at <http://wnet.isti.cnr.it/software/ttem> as free software, in order to make it possible for everyone to check, use and modify it.

III. MEASUREMENTS ENVIRONMENT

This paper compares simulation done using real traces of frame errors on a raw (before ARQ) Wi-Fi channel with synthetic traces of frame errors produced by commonly used models.

We performed a comprehensive measurement campaign [2] using laptops with standard Wi-Fi interfaces configured in ad hoc mode, at different fixed speeds of 1, 2, 5.5 and 11 Mb/s, fragmentation disabled, retransmission disabled, for different distances in different environments: open air, far from buildings, and office environment with varying numbers of thin walls between the laptops. We used custom software for sending the frames at precisely controlled time intervals and to receive them while registering the signal strength of received frames and the occurrences of lost frames, and we obtained traces of frame losses of various lengths and at different time resolutions.

The traces that are considered for this work are obtained in the authors' research institute ISTI in Pisa (IT), an office building. Fig. 1 depicts the arrangement of the relevant area of the building, at the first floor, where thin concrete walls delimit the rooms.

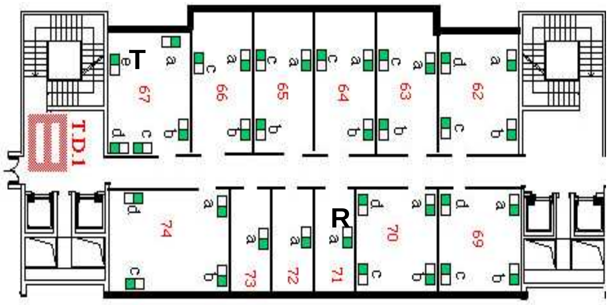

Fig. 1. Plan of the building where the measurements have been taken; T and R are the transmitter and receiver laptops used for the measurements.

The transmitter and receiver are two laptops equipped with a Debian GNU/Linux operating system. The transmitter sends a flow of 1000-byte unicast frames spaced by 5 ms, at a fixed 11 Mb/s rate with fragmentation and retransmission disabled. The receiver checks the sequence number inside the frames and keeps a trace of the lost ones. Since the channel occupancy of a frame is about 1.2 ms, this kind of measure traces the indoor channel conditions quite accurately.

IV. Commonly used packet error models

The most common model used to model frame error on a Wi-Fi channel is the Bernoullian process, that is, an error process where each frame experiences the same fixed probability of being corrupted, and frame losses are independent from each other [8], [9]. The main attractiveness of this model is its simplicity and its mathematical tractability. Only one parameter is needed to define it, usually the mean error probability p. A Bernoullian frame error model generates *burst* of consecutive frame errors (*error bursts*) whose length is geometrically distributed with mean equal to $1/(1\text{-}p)$. The distance between two bursts, that is, the *gap length*, is geometrically distributed as well, with mean equal to $1/p$.

A slightly more complex model defines the channel as being in one of two states, namely a *good* state where all frames are successfully received and a *bad* state where all frames are lost because of corruption [11], [12]. The mathematical modelling of this channel is done via a two-state Markov chain, usually referred to as a *Gilbert* or *Gilbert-Elliott model*. In fact, both the Gilbert model [13] and the Elliott model [14] are more complex, the first requiring three parameters and the second requiring four. This is why we refer to the simple good-bad model as a *bistable* model, where the probability of going from the good to the bad state is b and the probability of going from the bad to the good state is g. The stationary probabilities of being in the good and bad states are $P_g = g/(b+g)$ and $P_b = b/(b+g)$, respectively; P_b is also the mean frame error probability. The mean durations of the good and bad states are $1/b$ and $1/g$, respectively; this means that the channel exhibits error bursts of consecutive ones, whose mean length is $1/g$, which are separated by gaps of consecutive zeros whose mean length is $1/b$. Both burst and gap lengths have geometric distributions.

The models mentioned have no long-range memory: the Bernoullian process is memoryless, while the bistable process has memory limited to the previous state. As a consequence, both models have geometrically decaying correlations. However, the frame error traces we measured exhibit remarkably different statistics, which is the main ground for this investigation. As an example, we show a plot of the *error gap distribution*, a statistic originated with classical studies about bit-level models for telephone wires [13], and that may be significant in the context of TCP, where the length of error-free periods has a role as far as the goodput is concerned. The error gap distributions observed in the measured traces exhibit long polynomial tails, a peculiar trait that is not shared with commonly used models: both Bernoullian and bistable models produce error gap distributions with geometrically decaying tails, as depicted in Fig. 2 for the case of a Bernoulli synthetic trace.

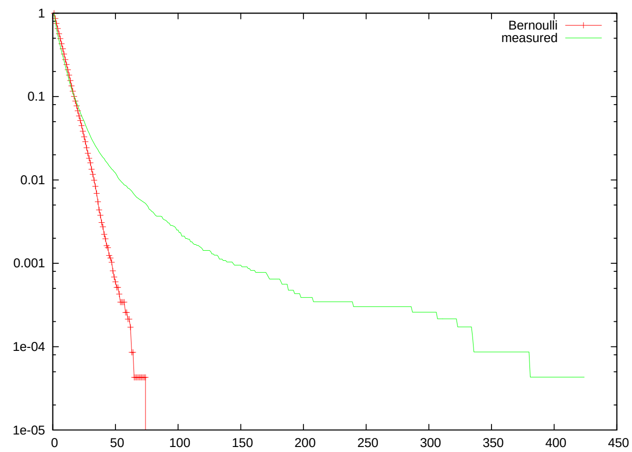

Fig. 2. Error gap length distributions for one observed frame error trace and a synthetically generated Bernoulli trace with the same average frame error rate.

V. TCP on synthetic versus measured packet loss

Even if the real traces exhibit significant statistical differences with respect to the Bernoulli and bistable models, choosing a model for simulation purposes need not be done on the basis of statistical similarity with the real process that is simulated, rather it should be done on the basis of how well the model behaves with respect to the intended purpose of the simulation. Therefore, it may well happen that a simple channel frame error model whose behaviour is apparently very different from the real packet error process is well suited for simulation in a given environment.

To this aim, we examined how well the two simple channel error models mentioned above perform, when used in modelling the behaviour of TCP in hybrid networks, by using the ns-2 simulator with the Ttem error model, which reads a trace of the channel state, applies the IEEE 802.11 ARQ algorithm and discards packets that cannot be received after a given number of retries, 7 in our case. The performance measure considered is the goodness of fit of a graph of *single-connection TCP goodput* versus mean frame error rate on the raw channel (before ARQ). We compare the goodput obtained with synthetic Bernoulli and bistable models to the goodput obtained using measured traces in the above-depicted indoor scenario.

First of all, we validate our simulation scenario by comparing Padhie's formula [15] with the results obtained with a synthetic Bernoulli trace. Padhie's formula quite accurately models the behaviour of a single TCP connection on a fixed-bandwidth link with uncorrelated segment losses due to either congestion or corruption, neglecting losses on the return channel and considering no link buffer. To compare the simulation's results with Padhie's formula, we run 10 instances of a single TCP connection for 2000 seconds, discarded the first 200 seconds to comply with Padhie's steady-state assumption, and considered the mean RTT, RTO and forward segment error rates of the 10 runs for frame error rates ranging from 2% to 30%. As shown in Fig. 3, the median values of the goodput over 10 ns-2 runs and Padhie's formula have a good agreement in the high error region, where the link buffer is mostly empty, while Padhie underestimates the goodput in the low error region, where the buffer influence becomes important.

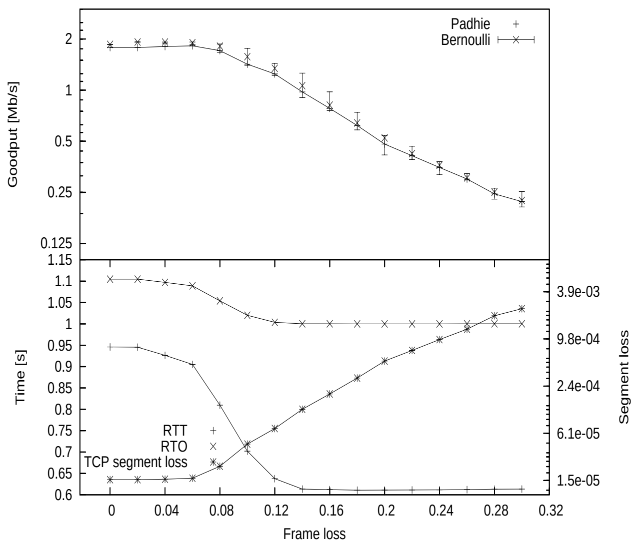

Fig. 3. Goodput as computed with Padhie's formula and with a trace with Bernoullian frame error.

We then made the comparison that is the aim of this paper. Taking the steady-state TCP single-connection goodput as a performance measure, we see how well different statistical models fit the measured frame error traces. We consider four different models in addition to the real traces, as shown in Table I.

Table I. Traces used in the simulation experiment.

(a)	Bernoulli process having the same frame error rate as the real traces
(b)	Bernoulli process having the same mean burst length as the real traces
(c)	Bernoulli process having the same mean gap length as the real traces
(d)	bistable process having the same frame error rate, mean burst length and mean gap length as the original traces
(e)	the real traces

Table I is worth some comments. When using a Bernoulli process — which is defined by a single parameter — fitting it to observed data means choosing a significant statistical parameter of the observed data. The choice of such a parameter is not obvious. In our case, the simplest choice of a parameter is the mean FER. However, there are some good reason why choosing a different parameter could be wiser. One candidate as an alternative to the FER is the mean burst length, an important parameter with respect to ARQ performance: the longer the error bursts, the higher the probability that ARQ cannot recover a lost frame. Another candidate is the mean gap length, which is related to TCP performance. As it is observed in [16], TCP performance is higher for higher burstiness of the segment error process, because the congestion window has a higher probability of becoming big if gaps are long. In fact, the measured error process has both longer bursts and longer gaps than a Bernoulli process with the same FER.

In Fig. 4, the goodput is plotted versus the frame error rate *of the real traces*.

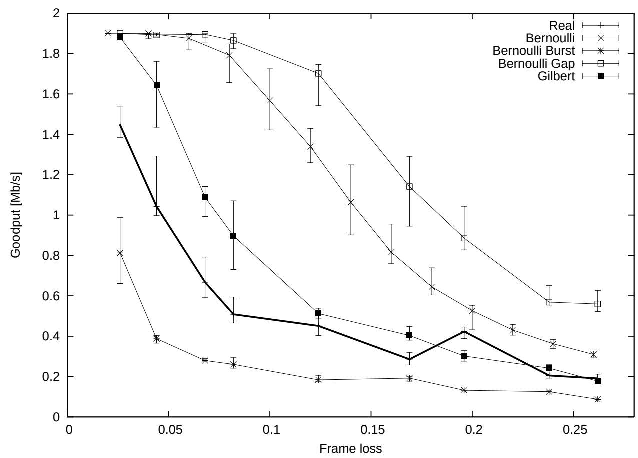

Fig. 4. TCP goodput versus frame error rate of real traces for different error models.

First of all, it is clear that (a), the simplest process, does not adequately model TCP's goodput at different frame error rates: the goodput with (a) is at least twice the one of (e); moreover, the behaviour is different, as (e) is convex, while (a) is not. Using the mean burst length as the parameter of interest gives the same behaviour as the real traces, as it can be seen by comparing (b) with (e), but the values are very different. As discussed above, this observation is a hint that the behaviour of ARQ is dominant in the chosen performance measure, that is, the TCP goodput. Apparently, the gap length, on the other hand, is much less important, as using it as the parameter for tuning the Bernoulli process does not yield satisfactory results: the trace for the (c) case is the one farthest from (e). One reason could be that, as shown in Fig. 2, the distribution of gap lengths in the real process has a polynomial, rather than exponential, tail implying that a Bernoulli process with the same mean gap length as the real traces has a lower FER than that of the real traces.

We considered until now a model with a single parameter. With the bistable model, we have two parameters, hence a greater flexibility. The three statistics of mean frame error rate, mean burst length and mean gap length are related, so they are defined by only two parameters, which implies that they can be simultaneously fitted using the bistable model. In fact, the bistable model fits the chosen performance measure best: the plot of (d) is nearest to (e), while still a considerable distance away.

VI. CONCLUSIONS

When choosing a model for the frame error rate process on a Wi-Fi channel, measurements coupled with simulation experiments using ns-2 show that none of the simplest models can be used to obtain an accurate estimate of the TCP goodput. Neither the Bernoulli process nor the good-bad bistable process with two parameters (sometimes referred to as Gilbert-Elliott process) are adequate when the main statistics of FER, mean burst length and mean gap length are considered.

This conclusion implies that more work is needed in this area, in order to provide researchers and simulation practitioners with reliable models for the frame error process at the channel level. The first direction of investigation should be directed to considering whether the simple Bernoulli or bistable processes could be used for TCP performance evaluation, by choosing parameters different from the FER and the mean burst or gap lengths. Should this turn out to be infeasible, the research should be directed to different, possibly more complex error models. Similar research methods should be used for estimating performance measure other than TCP's goodput, such as audio or video performance, ARQ performance or Wi-Fi channel capacity.

One thing that is conspicuously missing from the above analysis is the speed switching behaviour of the Wi-Fi interface. In order to model this behaviour, one would need channel measurements done at different speeds in different conditions in addition to choosing a speed switching algorithm, whose behaviour would greatly influence the performance measure under study, and that could itself be the object of research.

One more extension of this work could be considering the frame error process caused by atmospheric fading on the satellite channel, in addition to the one relative to the Wi-Fi channel.

REFERENCES

[1] IEEE Std 802.15.1-2002 IEEE Std 802.15.1 IEEE Standard for Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements Part 15.1: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks (WPANs).

[2] A. Annese, P. Barsocchi, N. Celandroni, and E. Ferro, “Performance evaluation of UDP multimedia traffic flows in satellite-WLAN integrated paths”, proceedings of 11th Ka Band Utilization Conference, Rome, September 2005.

[3] F. Davoli, E. Ferro, H. Mouftah, "Wireless access to the global Internet: mobile radio networks and satellite systems", guest editorial, Int. Journal of Communication Systems, vol. 16, no. 1, 2003.

[4] Gabriele Oligeri, “Misura e caratterizzazione di un canale Wi-Fi”, <http://fly.isti.cnr.it/didattica/tesi/Oligeri-Finotti/Oligeri.pdf>.

[5] M. Allman, S. Dawkinks, D. Glover, J. Griner, T. Henderson, J. Heidemann, S. Ostermann, K. Scott, J. Semke, J. Touch and D. Tran, “Ongoing TCP research related to satellites”, IETF, RFC 2760, February 2000.

[6] N. Celandroni and F. Potortì, “Maximising single connection TCP goodput by trading bandwidth for BER”, International Journal of Communication Systems, Vol. 16, No. 1, pp. 63-79, February 2003. Special issue on Wireless Access to the Global Internet: Mobile Radio Networks and Satellite Systems.

[7] D. Adami, S. Giordano, M. Pagano, and R. Secchi, “Modeling the Behavior of a DVB-RCS Satellite Network: an Empirical Validation”, proceedings of the International Working Conference ons Performance Modelling and Evaluation of Heterogeneous Networks, Ilkley (UK), July 2005.

[8] J.Jun, P.Peddabachagari, M.Sichitiu, Theoretical Maximum “Throughput of IEEE 802.11 and its Application”, Second IEEE International Symposium on Network Computing and Application (NCA'03), Cambidge (US), April 2003.

[9] Z. Tang, Z. Yang, J. He and Y. Liu, “Impact of bit errors on the performance of DCF for wireless LAN”, proceedings of the IEEE coonference on Communications, Circuits and Systems, vol. 1pp. 529-533, 2002.

[10] P. Chatzmisios, A. C. Boucouvalas and V. Vitsas, “Performance analysis of IEEE 802.11 DCF in presence of transmission errors”, proceedings of the IEEE International Conference on Communications (ICC 2004), vol. 7, pp. 3854-3858, June 2004.

[11] Michele Zorzi, Ramesh R. Rao, and Laurence B. Milstein, “On the accuracy of a first-order Markov model for data transmission on fading channels, *Proceedings of the International Conference on Universal Personal Communications*, pp. 211-215, November 1995.

[12] B. P. Crow, I. Widjaja, J. G. Kim and P.T. Sakai, “IEEE 802.11 wireless local area networks”, IEEE Communications Magazine, vol. 35, no. 9, pp. 116-126, 1997.

[13] E.N. Gilbert, “Capacity of a burst-noise channel”, *The Bell System Technical Journal*, No. 39, pp. 1253-1265 September 1960.

[14] E. O. Elliott, “Estimates of error rates for codes on bursty-noise channels”, Bell System Tech. J., 42(9):1977, September 1963.

[15] J. Padhie, V. Firoiu, D.F. Twsley, and J.F. Kurose, “Modeling TCP Reno performance: a simple model and its empirical validation”, IEEE/ACM Trans. Networking, vol. 8, pp. 133-145, April 2000.

[16] E. Altman, K. Avrachenkov, C. Barakat, “A stochastic model of TCP/IP with stationary random losses”, Computer Communication Review, Vol.30, No.4, October 2000, pp. 231-242.