Modelling TCP performance in mobile DVB-T receivers

GEORGIOS GARDIKIS
Mobile Radiocommunications Laboratory
National Technical University of Athens
9, Iroon Polytechneiou, Zografou, Athens
GREECE
gardikis@iit.demokritos.gr

ANASTASSIOS KOURTIS
Institute of Informatics and Telecommunications
NCSR "Demokritos"
Patriarchou Gregoriou, Agia Paraskevi, Athens
GREECE

PHILIP CONSTANTINOU
Mobile Radiocommunications Laboratory
National Technical University of Athens
9, Iroon Polytechneiou, Zografou, Athens
GREECE

Abstract: Contemporary DVB-T networks, when combined with wireless data links for the interaction data, can provide an effective and affordable solution for Internet and data service provision to portable and mobile end users. This paper presents the performance of TCP in a hybrid DVB-T/GPRS network under a fading environment in the downlink. A mathematical model is derived to approximate TCP throughput based on the parameters of mobile reception (speed, SNR margin) in the case of a flat-fading Rayleigh channel.

Key-Words: Interactive digital terrestrial television, TCP, Flat Rayleigh channels

1 Introduction

The performance and error tolerance of DVB-T in fading environments has been extensively investigated via simulations and field trials. However, most research efforts in this area focus in the case of digital television program transmission, where BER requirements are very strict – since a single corrupt bit can cause annoying picture impairments. Certain SNR thresholds exist for every transmission mode and fading model, below which the quality of service is unacceptable. This is not the case, though, with error-resilient data transfer protocols like TCP, where occasional data loss can be compensated. This paper analyzes the behaviour and performance of TCP within a fully functional DVB-T/GPRS hybrid system under a fading downlink environment. An analytical approximation of the measured throughput values is derived, leading to a generalized mathematical model describing the performance of TCP in a DVB-T receiver suffering from flat Rayleigh fadings.

The network model examined is depicted in Fig.1. In this case, TCP communication follows an asymmetric path, with requests and acknowledgments conveyed by the GPRS uplink, while data destined for the end user are routed via the DVB-T downlink.

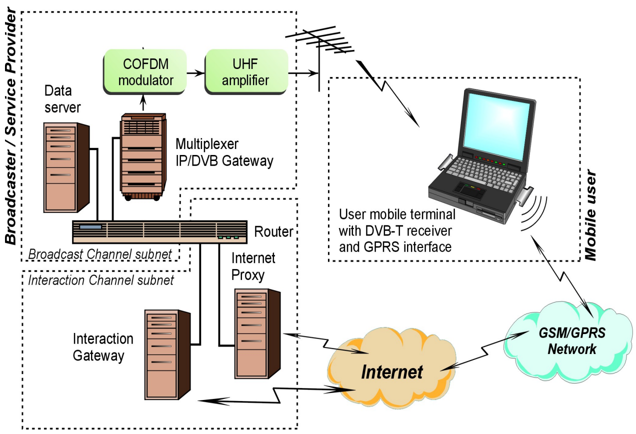

Fig.1. Overall architecture of the hybrid GPRS/DVB-T testbed

2 Approximation of the flat-fading DVB-T channel with a two-state stochastic model

Assuming the GPRS uplink reliable and error-free, this paper will investigate the effect of the mobile propagation channel whose fluctuations can cause unrecoverable errors in the receiver, resulting in packet loss in the TCP data flow[1].

It is known that, due to the "cliff effect" which is clearly observable in DVB-T reception, the downlink receiver can substantially be found in two discrete states at any specific moment: either the error-free mode ("Good" state – when the block and Viterbi decoders manage to correct the errors introduced by the channel) or the totally erroneous mode ("Bad" state – when the decoding process fails and the baseband signal is impossible to be reconstructed). The transition between the two states occurs at an SNR margin of 1 dB at most[2]. The SNR threshold value under which the receiver falls in the "Bad" state depends on the transmission mode and can be found in the Annex A of the DVB-T specification (ETS 300 744).

Making the assumption that the DVB-T signal traverses a flat-fading Rayleigh channel, this transition sequence between the "Good" and "Bad" states can be approached by a two-state first-order Markov chain ("On-off" or Gilbert-Elliot model) [4][5]. The system resides in each state for duration τ_G, τ_B respectively. The PDFs of these time intervals are exponential with means $\overline{\tau_G}, \overline{\tau_B}$.

Following this assumption, the wireless DVB-T link of the network testbed can be replaced by a direct wired route which goes through an electronic RF attenuator, introducing a 40dB drop in the traversing signal at the activation of a TTL pulse. This attenuator is controlled by a dedicated computer running a simulation software module developed for the needs of this paper. The software activates and deactivates the attenuator at random exponential durations τ_B, τ_G respectively, with user-defined means $\overline{\tau_B}, \overline{\tau_G}$. The activation of the attenuator corresponds to the case of the received SNR being below the threshold (the distance in dB between the actual SNR and the threshold value is of no importance, since the decoding process has failed), and forces the receiver to drop to the "Bad" state. The SNR rise above the threshold (attenuator deactivation) resumes the receiver to the "Good" state.

Although Internet connectivity is supported by this configuration (via the Internet proxy in Fig.1), the behaviour and performance of TCP was tested through a data transfer session from the Linux-based local data server to the Client PC (WinXP) – to ensure that throughput measurements are not affected by uncontrolled bottlenecks outside the local network. The data transfer session was generated by the Iperf network benchmarking suite and lasted 2 minutes each. Series of throughput measurements were made for different receiver TCP window sizes ranging from 8K to 512K and for various combinations of $\overline{\tau_B}, \overline{\tau_G}$. The results for window sizes of 16K and 128K are displayed in Fig.2a,2b and are irrelevant of the DVB-T transmission mode. The absence of time interleaving in the DVB-T system makes it impossible for the receiver to recover from a flat deep fade, regardless of the code rate and the modulation used. Of course, in the real world the transmission mode does matter, because the SNR threshold changes and a several-dB fade will affect the more robust modes less, as the receiver will spend less time in the "Bad" state.

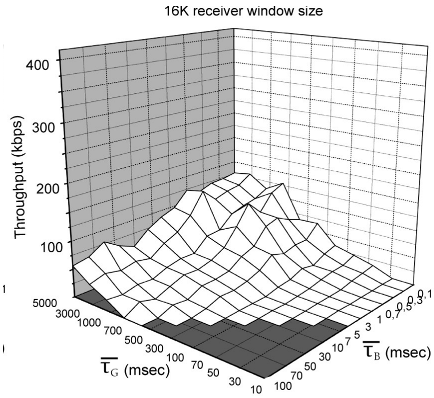

Fig.2a. Measured TCP throughput vs. average state duration, 16K window

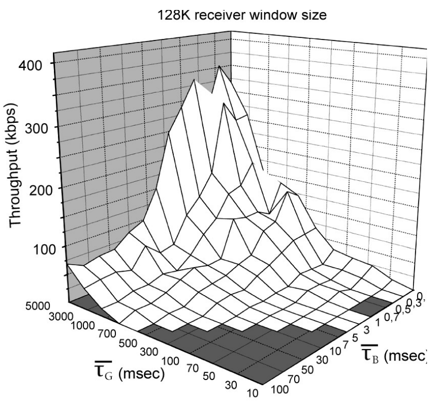

Fig.2b. Measured TCP throughput vs. average state duration, 128K window

As depicted in Fig.2, with 128K size of TCP window, higher throughput values are obtained under light fading conditions as the increased receive buffer compensates for the high round-trip delay (about 900 msec) inserted mainly by the GPRS Interaction Channel. Larger window sizes do not affect performance, while smaller ones do not have better behaviour in the case of frequent packet loss. With TCP Selective Acknowledgements (SACKs) option enabled at both the receiver and the sender, TCP can efficiently and quickly recover from packet losses regardless of the size of the window. Therefore, the value of 128K can be selected as optimal in this type of network.

3 Modelling TCP throughput for mobile receivers at specific SNR

The specification of TCP throughput under a strictly defined average "Good" and "Bad" duration can help in deriving a mathematical model which links TCP performance with the speed and the SNR value at the receiver.

The 128K graph at Figure 2, drawn in almost logarithmic axes, can be approached by a log-binomial function whose coefficients are defined via a computer-aided curve fitting procedure:

$$T(\overline{\tau_G},\overline{\tau_B}) = Pos\left|\left(-1.753\cdot\ln(\overline{\tau_G}) + \ln(\overline{\tau_B}) - 2.315\right)^2 - 87.088\right| \quad (1)$$

where Pos[f] implies the positive part of f, as the binomial can also reach negative values. These negative values actually correspond to zero-throughput points. T in (1) is a good approximation to the experimental results, presenting a standard deviation of 22kbps when compared with the measured values.

Expression (1) can be used as an intermediate step to link TCP throughput with real-life measurable quantities, like terminal relative speed and received SNR. Mean durations $\overline{\tau_B},\overline{\tau_G}$ can be expressed as functions of receiver speed and SNR margin between the RMS received value and the threshold. For Rayleigh flat fading channels, $\overline{\tau_B}$ is equal to the average fade duration (AFD):

$$\overline{\tau_B} = \frac{e^{\rho^2}-1}{\rho f_m\sqrt{2\pi}} \text{ and } \overline{\tau_G} = \frac{1}{N_R} - \overline{\tau_B}\text{, where}$$
$$N_R = \sqrt{\frac{\pi}{\sigma^2}}\,\rho\; f_m\, e^{-\rho^2} \quad (2)$$

N_R is the positive level crossing rate (LCR – the rate at which the receiver visits the "Good" state), $f_m = \frac{v}{\lambda} = \frac{vf}{c}$ the maximum Doppler frequency shift (f=602MHz, UHF channel 37) and ρ the ratio of the threshold over the received signal, which can be expressed as a SNR ratio. By combining (2) and (1), the final formula is derived:

$$T(M,v) = Pos\left[-87.088 + \left(\ln\left(e^{10^{-\frac{M}{10}}} - 1\right) + 0.753\ln v - 0.086M - 7.23\right)^2\right] \quad \textbf{(3)}$$

Expression (3) approximates TCP bulk throughput in a hybrid DVB-T/GPRS network (window size set to 128K minimum) where v is the speed of the receiver (in km/h) and M the distance in dB between the received SNR and the threshold SNR for the specific DVB-T transmission mode. The graph of T(M,v) is shown in Fig.3.

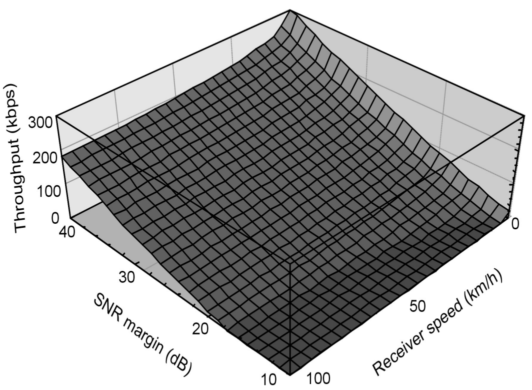

Fig.3. Analytical approximation of TCP throughput vs. speed and SNR margin

From the graph it is clear that TCP performance is significantly boosted in slow-moving receivers and even more in almost stationary conditions. This happens because in slower speeds the average duration between two successive deep fades is increased and TCP has the time to increase its congestion window by sending more data to the network. From the graph it is also deduced that a typical receiver speed in the order of 20 km/h corresponds to a need for a SNR margin of at least 20dB to achieve acceptable performance. In the case of adequate SNR margin, then the DVB-T/GPRS combination can offer a real TCP throughput of 200kbps or more, allowing for Internet access, Web browsing and file transfer at an acceptable rate.

4 Conclusion

This paper presented the results of measurements on the performance of TCP over a hybrid asymmetric GPRS/DVB-T Network. The measurements show even that a narrowband access technology like GPRS, which can be utilised at a very low cost, when combined with the robust DVB-T broadcast downlink, can result in download rates equivalent to those of modern wired access technologies. Data transfer measurements in simulated fading conditions were presented and a mathematical model was derived to show that ADSL-equivalent and even higher bit rates can be achieved in a fading environment without need for perfect reception and high SNR values. A viable and affordable solution for high-speed Internet access to mobile and portable end users is thus provided.

References:

[1] Xilouris, G., Gardikis, G., Pallis, E., and Kourtis, A., Reverse Path Technologies in Interactive DVB-T Broadcasting, *Proc. IST Mobile and Wireless Telecommunications Summit*, Thessaloniki, Greece, June 2002, pp. 292-295

[2] Weck, C., VALIDATE Field Trials of Digital Terrestrial Television (DVB-T), *Proc. NAB '98*

[3] Stare, E., Mobile reception of 2K and 8K DVB-T signals, *Proc. IBC '98*, pp.473-478

[4] Wang, H. and Moayeri, N., Finite-State Markov Model - A Useful Model for Radio Communication Channels, *IEEE Trans. on Vehicular Technology*, February 1995

[5] Zorzi, M., Rao, R. and Milstein, L. On the

accuracy of a first-order Markov model for data transmission on fading channels, *Proc. ICUPC'95*, Tokyo, Japan