MODELLING AND CONTROL OF IP TRANSPORT IN CELLULAR RADIO LINKS [1]

N. Möller ** **C. Fischione** * **K. H. Johansson** **
F. Santucci * **F. Graziosi** *

* *University of L'Aquila, Dpt. of Electrical Engineering*
Poggio di Roio, I-67040 L'Aquila, Italy
e-mail: {fischione, santucci, graziosi}@ing.univaq.it
** *Royal Institute of Technology*
Dept. Signals, Sensors and Systems
Automatic Control, SE-100 44 Stockholm, Sweden
e-mail: {niels, kallej}@s3.kth.se

Abstract: A fundamental assumption of the TCP protocol is that packet losses indicate congestion on the network. This is a problem when using TCP over wireless links, because a noisy radio transmission may erroneously indicate congestion and thereby reduce the TCP sending rate. Two partial solutions, which improve the quality of the radio link, are power control and link-layer retransmissions. We consider a radio channel with multiple users and traffic classes, and investigate how parameters in the radio model influences TCP-related quality measures, such as the average delay and the probability of spurious timeout and spurious fast retransmit. The results indicate that the outer loop power control is robust to uncertainties in the radio model. This robustness property supports separation between the radio layer design and the IP and TCP layers.

Keywords: TCP, WCDMA, Power control, Feedback control, ARQ

1. INTRODUCTION

Third generation and up-coming mobile radio systems, such as those based on the WCDMA radio interface, are developed to enable, among others, wireless access to the Internet. TCP/IP is the protocol suite having the widest diffusion for the transport infrastructure of wired networks, and therefore we expect TCP to play an increasingly important rôle in hybrid wired/wireless networks.

Poor TCP performance over wireless links is a well-known problem. The traditional explanation for poor TCP performance is that the wireless link drops packets due to noise on the radio channel, and that TCP interprets all packet losses as indications of network congestion. The performance of a radio link transmission can be improved with link-layer retransmissions schemes, which are common in cellular systems. The link-layer retransmission scheme transforms a lossy link, with fairly constant delay, into a link with few losses but random delays.

Predicting TCP performance degradation over wireless links and devising possible improvement strategies have been a subject of research in recent years (see e.g. (Xylomenos *et al.*, 2001)). In fact, wireless communication systems are prone to channel impairment and interference, and typical end-to-end congestion and error control mecha-

[1] This work was partially supported by the European Commission through the projects EuroNGI and HYCON, and by the Swedish Research Council.

nisms of TCP (Xylomenos *et al.*, 2001), (Barman *et al.*, 2004) basically fail in interpreting the sequence of events produced by lower layers in the wireless protocol architecture. Among improvement strategies, some contributions (Mascolo *et al.*, 2001; Ludwig and Katz, 2000) have been oriented to propose TCP variants over fairly general classes of wireless links. However, this approach has the significant drawback that TCP protocols currently implemented in communicating terminals should be replaced. An alternative approach is instead focused on exploring solutions for the improvement of the wireless interface, with particular emphasis on UMTS-related contexts.

Nevertheless, relevant efforts have been devoted to model the "TCP over wireless" scenario, as a fundamental preliminary step for devising suitable solutions. Specifically, in (Liu *et al.*, 2002) forward error correction coding to add redundancy to the packets is addressed. More detailed radio link models with no retransmissions are taken into account in (Abouzeid *et al.*, 2000). TCP behavior over wireless links with frame errors and retransmissions described by a two-state Markov process is investigated in (Pan *et al.*, 2002). The same model is used in (Canton and Chahed, 2001) and (Chahed *et al.*, 2003), where a simple scheme of power control is also included. In particular, the use of two-state Markov chains to model the radio channel has been extensively studied in last years (see also (Rossi *et al.*, 2004)). In (Möller and Johansson, 2003), a Markov chain with multiple states is used to model a power-controlled channel, and derive TCP performance properties. In (Khan *et al.*, 2000), TCP throughput over a link is simulated, for various radio channel conditions and link-layer retransmission schemes. In (Hossain *et al.*, 2004), the investigation has concerned the forward link power allocation and rate adaptation for TCP throughput maximization in a WCDMA wireless system, where perfect channel estimation is assumed.

In this paper, we show that for one common type of power control, the essential TCP/IP behavior seems to be independent of radio parameters. The power control seems robust enough to hide uncertainty in the radio model. To be more precise, we show that with a simple SINR-based outer loop in the power control, changes to the performance of the power control inner loop has virtually no effect on TCP-layer properties. This robustness property supports a separation between the radio layer and the IP and TCP layers.

To explain this observation, the channel is described by the loss curve $f(r), \text{FER} = f(\text{SINR})$, which gives the expected frame error rate for a given signal to interference and noise ratio. The objective of the outer loop power control is to keep

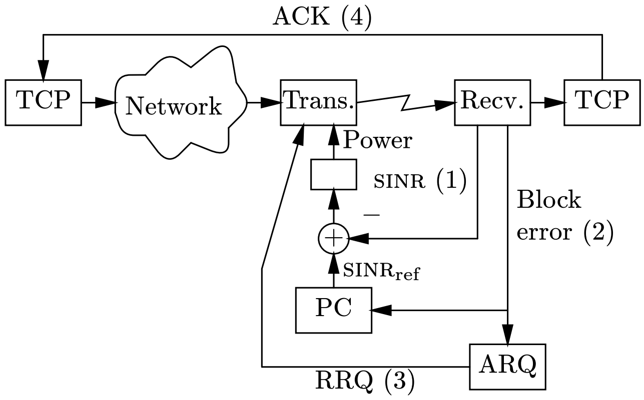

Fig. 1. System overview. Four feedback loops: Inner loop power control (1), outer loop power control (2), link-layer retransmissions (3), and end-to-end congestion control (4).

the FER and SINR close to the desired operating point on this curve. From this point of view, it is natural that it is the shape of the curve close to the operating point that is relevant for layers above power control.

The rest of this paper is organized as follows. In Section 2 we describe our models for power control and link-layer retransmissions. In Section 3 we explain how to use these models to derive TCP/IP-related properties. In Section 4, we introduce a more complex radio channel that takes multi-access interference into account. Section 5 describes and explains the results we get for this radio model.

2. SYSTEM MODEL

When using TCP over a wireless link, there are several interacting control systems stacked on top of each other, as illustrated in Figures 1. At the lowest level, the transmission power is controlled in order to keep the signal to noise and interference ratio (SINR) at a desired level. This is a fast inner loop intended to reject disturbances in the form varying radio conditions. On top of this, we have an outer power control loop that tries to keep the frame error rate constant, by adjusting the target SINR of the inner loop. Next, we have local, link-layer, retransmissions of damaged frames. Finally, we have the end-to-end congestion control of TCP.

We will describe these layers in turn. In Section 4 we will describe a more detailed radio model that takes multi-access interference into account.

2.1 Power control

The typical SINR-based power control uses an inner loop that tries to keep the signal to interference ratio (SINR) close to a reference value

SINR$_{\text{ref}}$. This loop often has a sample frequency of 1500Hz, and a one bit feedback that is subject to a delay of two samples, i.e., 1.3 ms. In simulations (Gunnarsson and Gustafsson, 2002), the inner loop is able to track the reference SINR$_{\text{ref}}$ within 2-3 dB, with a residual oscillation due to the delay and the severe quantization of the power control commands in the inner loop. The period of this oscillation is typically less than 5 samples, i.e., 3.3 ms.

As there is no simple and universal relationship between the SINR and the quality of the radio connection, there is also an outer loop that adjusts SINR$_{\text{ref}}$. This loop uses feedback from the decoding process; in this article we assume that the power control outer loop is based on frame errors.

As it is hard to estimate the frame error rate accurately, in particular if the desired error rate is small, one approach is to increase SINR$_{\text{ref}}$ significantly when an error is detected, and decrease the SINR$_{\text{ref}}$ slightly for each frame that is received successfully. It is interesting to note that this strategy resembles the TCP "additive increase, multiplicative decrease" congestion control strategy. We will discuss this approach in more detail later. For a survey of modern power control techniques for systems such as WCDMA, see (Gunnarsson and Gustafsson, 2002).

The goal of the power control is to keep the frame error rate (FER) close to a given value p. For data traffic in UMTS, $p = 0.1$ is a common choice, and that is what we will use. The chosen frame error rate is a deployment trade-off, between transmission quality and the required number of base stations.

2.1.1. Markov model The outer loop of the power control sets the reference value for the SINR. Given a particular reference value SINR$_{\text{ref}} = r$, the obtained SINR is a stochastic process. Together with the coding scheme for the channel, we get an expected probability for frame errors. If the coding scheme is fixed, the probability of frame errors is given by a function $f(r)$.

The outer loop of the power control uses discrete values for SINR$_{\text{ref}}$. One way to keep the frame error probability close to the desired probability p is to change SINR$_{\text{ref}}$ by fixed steps, based on a step size Δ. Whenever a radio frame is received successfully, SINR$_{\text{ref}}$ for the next frame is decreased by Δ. And whenever a radio frame is damaged, SINR$_{\text{ref}}$ for the next frame is increased by $K\Delta$, where $1/(1+K) = p$. The value of Δ is an important control parameter, which determines the performance of the power control.

For an integer K, the varying SINR$_{\text{ref}}$ can be viewed as a discrete Markov chain (Sampath *et al.*, 1997). To do this, we have to make the assumption that frame errors depend only on the value of SINR$_{\text{ref}}$ when the frame is transmitted; there is no relevant history except the SINR$_{\text{ref}}$. In our case $p = 0.1$ implies $K = 9$.

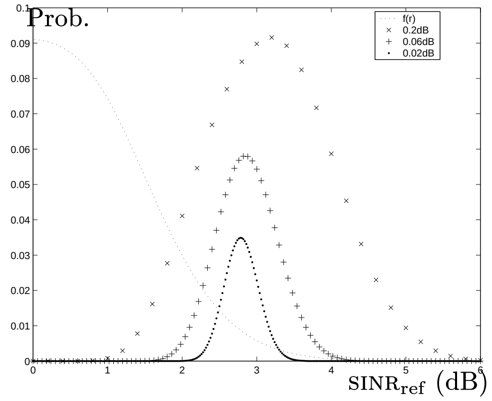

Fig. 2. Stationary distribution for the power control. Each mark represents one state of the power control, the corresponding value of SINR$_{\text{ref}}$, and its stationary probability π_i. The dotted curve is the threshold-shaped function $f(r)$, scaled to fit in the figure, which represents the frame error probability as a function of SINR$_{\text{ref}}$.

We get a finite Markov chain by truncating it at the ends where $f(r) \to 1$ (all frames are lost) and $f(r) \to 0$ (no frames are lost), and it is straight forward to compute the stationary distribution, denoted π_k.

Figure 2 shows the threshold function $f(r)$, together with the stationary distribution for three values of Δ. The function $f(r)$ corresponds to a BPSK channel (Möller and Johansson, 2003).

Consider the dependency on the step size Δ. When the step size is decreased, the stationary distribution gets more narrow, which means that the received SINR, and hence FER, becomes more and more concentrated. In the limit, we will get a loss process where each frame is lost with probability $p = 0.1$ independently of other frames.

To examine what happens for larger Δ, divide the states into three subsets. The *loss set*, consisting of states with FER > 0.9, the *success set*, consisting of states with FER < 0.0001, and the *middle set*, consisting of the rest of the states. When Δ is increased, fewer and fewer states will belong to the middle set. In the limit, there will be at most one state that is in the middle set. Then the power control will stay in the success set most of the time. It will move one or very few steps below the success set into the loss set, generate one isolated loss, then make a large step back into the success set. The state will stay for about 9–10 time slots in the success set before it next enters the loss set. This implies that the loss process is almost periodic.

When focusing on power control performance, there are other important qualities that are influenced by the choice of Δ. A small Δ gives a longer system response time, while a large Δ will result in larger average transmission power, which limits both the battery life of devices, and the system capacity in terms of the maximum number of simultaneous users that can be accommodated without interference collapse of the system.

2.2 Link-layer retransmissions

The simplest way to transmit IP packets over the wireless link is to split each IP packet into the appropriate number of radio frames, and drop any IP packet where any of the corresponding radio frames were damaged. But as is well-known, TCP interprets all packet drops as network congestion, and its performance is therefore very sensitive to non-congestion packet drops. An IP packet loss probability on the order of 10% would be highly detrimental.

There are several approaches to recover reasonable TCP performance over wireless links. In this paper we concentrate on a local and practical mechanism: The link can detect frame damage through error correction codes (error detection is needed for power control anyway), and it can use that information to request that damaged frames be retransmitted over the link. This capability is an option in standard wireless network protocols, see (Bai *et al.*, 2000) for an evaluation of these options in the IS-2000 and IS-707 RLP standards.

The effect of link-layer retransmissions is to transform a link with constant delay and random losses into a link with random delay and almost no losses. Alternative approaches include changes to the TCP algorithms, e.g. Eifel (Ludwig and Katz, 2000), and TCP Westwood (Mascolo *et al.*, 2001), and the use of forward error correction coding to add redundancy to the packets, either end-to-end as in (Lundquist and Karlsson, 2002) or at the link as in (Liu *et al.*, 2002).

There are several schemes for link level retransmission. We will consider one of the simpler, the (1,1,1,1,1)-Negative Acknowledgement scheme (Khan *et al.*, 2000), which means that we have five "rounds", and in each round we send a single retransmission request. When the receiver detects that the radio frame in time slot k is damaged, it sends a retransmission request to the sender. The frame will be scheduled for retransmission in slot $k+3$ (where the delay 3 is called the RLP NAK guard time). If also the retransmission results in a damaged frame, a new retransmission

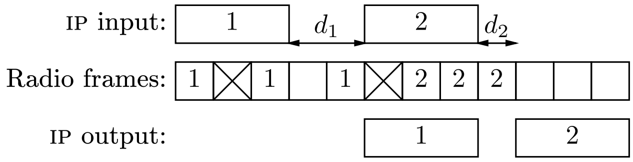

Fig. 3. Overlaying IP transmission on top of the frame loss process and frame retransmission scheduling. Two IP packets, corresponding to $n = 3$ radio frames each, are transmitted over the radio link. The second and sixth frames are damaged, crossed out in the figure, and scheduled for retransmission three frames later. The resulting retransmission delay for the i:th IP packet is denoted d_i.

request is sent and the frame is scheduled for retransmission in slot $k + 6$. This goes on for a maximum of five retransmissions.

In the next section, we put together the frame loss process and the retransmission scheduling, to derive IP level properties of the link.

3. TCP/IP PROPERTIES

When transmitting variable size IP packets over the link, each packet is first divided into fix size radio frames. We let n denote the number of radio frames needed for the packet size of interest. For the links we consider, we have $n \leq 10$.

We can overlay the IP packets on top of the frame sequence and simulate the retransmission scheduling, as illustrated in Figure 3.

Consider the system at a randomly chosen start time, with the state of the power control distributed according to the stationary distribution. For any finite loss/success sequence (for example, the second and the sixth block damaged, the rest received successfully), we can calculate the probability by conditioning on the initial power control state and following the corresponding transitions of the Markov chain. We can then use these probabilities to investigate the experience of IP packets traversing the link.

In the rest of this section, we explain how to use these loss/success-sequence probabilities to derive the average values of TCP/IP properties of interest.

3.1 Delay

As a link employing link-layer retransmission yields a very small packet loss probability, the most important characteristic of the link is the packet delay distribution. This distribution can be computed explicitly from the models described above, including its average.

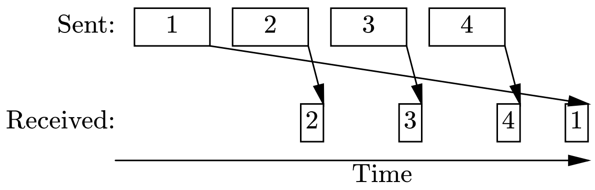

Fig. 4. Reorder triggers fast retransmit

3.2 Reorder—fast retransmit

If a packet in a TCP stream is delayed so much that it lags behind three other packets (see Figure 4), the acknowledgements for the next three packets are "duplicated", and the sender will go into fast retransmit/fast recovery mode. That means that from the point of view of TCP and its congestion control, the packet is lost.

Computing the precise number of fast retransmit events is not trivial. However, we can estimate the probability that fast retransmits occur as follows. Consider a randomly chosen time, when the power control state is governed by the stationary distribution π_i. Assume that four IP packets, each of size n radio frames, arrive to the radio link back-to-back. We then consider all possible loss/success sequences, and sum the probabilities of the sequences which cause the first packet to be delayed long enough to be the *last* of the four packets to be received successfully and completely. This gives us the probability of spurious fast retransmit, denoted $\mathrm{P_{FR}}$.

3.3 Spurious timeout

A timeout event occurs when a packet, or its acknowledgement, is delayed too long. Let RTT_k denote the round trip time experienced by packet k and its corresponding acknowledgement. The TCP algorithms estimates the mean and deviation of the round trip time. Let $\widehat{\mathrm{RTT}}_k$ and $\hat{\sigma}_k$ denote the estimated round trip time and deviation, based on measurements up to RTT_k. TCP then computes the timeout value for the next packet as $\widehat{\mathrm{RTT}}_k + 4\hat{\sigma}_k$ (Jacobson, 1988) which means that the probability that packet k causes a timeout is given by

$$\mathrm{P_{TO}} = P(\mathrm{RTT}_k > \widehat{\mathrm{RTT}}_{k-1} + 4\hat{\sigma}_{k-1})$$

We assume that the values RTT_k are identically and independently distributed, according to the delay distribution computed as in in Section 3.1. For simplicity, we also assume that the estimates $\widehat{\mathrm{RTT}}_k$ and $\hat{\sigma}_k$ are perfect and equal to the true mean and standard deviation of RTT_k.

4. MULTI-ACCESS INTERFERENCE

In this section, we model the physical layer for the reverse link of a single-cell asynchronous BPSK DS/CDMA system, incorporating the behavior of the abovementioned power control mechanisms. $S = 3$ classes of mobile users are considered, where each class is associated to a traffic source type (data, video, voice, ...) and each user to a mobile device.

The generic Class i is characterized by its own bit rate R_i (or the bit interval $T_i = \frac{1}{R_i}$) and SINR average requirement (average target SINR for the outer loop power control). The same fixed bandwidth W, and thus the same chip interval T_c, is allocated to every user. Therefore, each class of users has a processing gain $G_i = \frac{W}{R_i}$. There are K_i active users of the Class i (with $i = 0, ..., S-1$). Following the same approach outlined in (Santucci *et al.*, 2003) and (Fischione *et al.*, 2002), it can be shown that the SINR at the output of a coherent correlation receiver matched to signal related to the generic user 0 of Class 0 (user of interest involved in the TCP connection), has the following expression:

$$\mathrm{SINR}(P(t), \xi(t), \nu(t)) = L^{-\frac{1}{2}}(t) \qquad (1)$$

where

$$\begin{aligned} L(t) &= De^{-\xi_{00}(t)} \\ &+ \sum_{k=1}^{K_0-1} Ae^{\xi_{0k}(t)-\xi_{00}(t)}\nu_{0k}(t) \\ &+ \sum_{j=1}^{S-1}\sum_{k=0}^{K_j-1} B_j e^{\xi_{jk}(t)-\xi_{00}(t)}\nu_{jk}(t) \end{aligned}$$

and $A = \frac{1}{3G_0}$, $B_j = \frac{P_j(t)}{3P_0(t)G_0}$ and $D = \frac{N_0}{2P_0(t)T_0}$. The meaning of various variables and parameters in the above expressions is as follows:

- $N_0/2$ denotes the two-sided power spectral density of thermal noise at the receiver input;
- $P_i(t)$ denotes the power level at the receiver input for the generic user signal of Class i, and depends on longer term updates induced by the outer loop power control mechanism;
- $\nu_{ik}(t)$ is a binary random process indicating the activity status (On/Off) of the source at time t; its first order probability mass function is such that $\mathrm{P}(\nu_{ik} = 1) = \alpha_i$ and $\mathrm{P}(\nu_{ik} = 0) = 1 - \alpha_i$, where α_i is said to be the activity factor of sources of the Class i;
- $\xi_{ik}(t)$ denotes the residual (inner loop) power control error for the user signal k of Class i and is represented (in log units) by a zero mean Gaussian process with standard deviation σ_i (in dB).

We assume independence is assumed between any pair of the above processes.

In order to simplify computation, we assume a typical log-normal model for $P_i(t)$, which leads to a log-normal approximation for L (Fischione *et al.*, 2004),

$$L(t) \approx e^{Z(t)} , \qquad (2)$$

where the first and second moments of the Gaussian Z can be expressed in terms of radio model parameters, and the first and seconds moments of the Gaussian variables $\log P_i$.

In the simulations, we will focus on the parameter σ_0, the standard deviation of ξ_{0k}, the control error for the inner loop power control for users of class 0.

To apply our Markov chain machinery, we need the function f, with $\text{FER} = f(\text{SINR}_{\text{ref}})$. We get this from the expected bit error rate

$$\text{BER}(\text{SINR}_{\text{ref}}) = \text{E}(Q(\text{SINR}(P(t), \xi(t), \nu(t)))) \quad (3)$$

where the expected value is taken by considering P_0, the transmission power for users of class 0, kept fixed and selected as to get precisely SINR_{ref}, and then take the expected value over the other stochastic variables.

The frame error rate $\text{FER}(\text{SINR}_{\text{ref}})$ is then expressed in terms of the bit error rate and the error correction parameters.

5. NUMERICAL RESULTS

We now presents simulation results, where we vary σ_0, the variance of the control error of the inner power control loop, and study the effects on our measures mean delay, spurious timeout probability (P_{TO}) and spurious fast retransmit probability (P_{FR}).

The parameters used in the simulation are as follows: Three traffic classes, from the UMTS specification (3GP, n.d.). The number of users in the three classes are $K_0 = 20, K_1 = 5$, and $K_2 = 50$, activity factors are $\alpha_0 = 0.2, \alpha_1 = 0.4$, and $\alpha_2 = 0.7$, representative for interactive data traffic, low rate video, and voice. We focus on class 0 which is used for TCP traffic. The channel bit rate is 240 Kbit/s. The transmission time interval is 10 ms, corresponding to 2400 bits per frame. The number of correctable bit for each frame is set to $C = 8$. The Gaussian $\log P_1, \log P_2$ are assumed to have mean 2 dB and standard deviation 6 dB. We also set $N_0/T_0 = -13$ dB. See (Fischione *et al.*, 2005) for further details.

For the variance of the control error for the inner loop, we set $\sigma_1 = \sigma_2 = 1$ dB, and vary σ_0. The power control step size is set to $\Delta = 0.02$ dB.

In Figure 5, we see that the average delay is not influenced by σ_0. Since a larger packet implies more opportunities for frame errors, and retransmissions, the average delay increases with n.

In Figure 6, we see that the spurious timeout probability increases slightly with increased packet size and with increased σ_0. The dip for $\sigma_0 = 1.5$ may be due to the approximation deteriorating for large σ_0.

In Figure 7, we see that the probability for severe reorder, leading to spurious fast retransmit, is significant for $n = 1$, much smaller for $n = 2$ and negligible for $n \geq 3$. Again, there's only a weak dependence on σ_0.

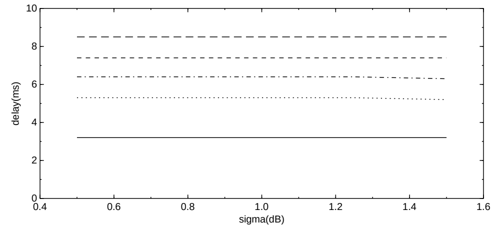

Fig. 5. Average delay, as a function of σ_0. The curves correspond to different packet sizes: $n = 1, \ldots, n = 5$, drawn bottom to top.

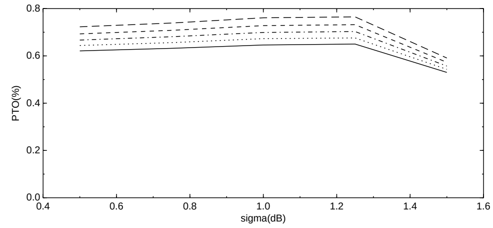

Fig. 6. Spurious timeout probability, as a function of σ_0. The curves correspond to different packet sizes

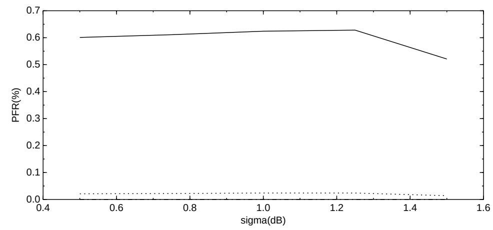

Fig. 7. Spurious fast retransmit probability, as a function of σ_0. The curves correspond to different $n = 1$ (solid) and $n = 2$ (dotted). We get $\text{P}_{\text{FR}} = 0$ for $n \geq 3$

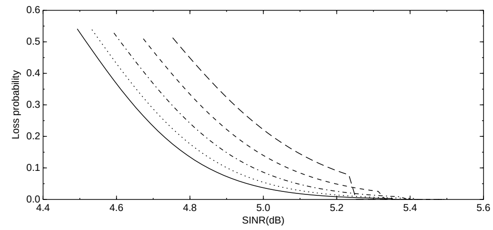

Fig. 8. Loss probability as a function of SINR_{ref}. For $\sigma_0 = 0.5, 0.75, 1.0, 1.25, 1.5$, drawn bottom to top

To explain this independence of σ_0, consider Figure 8, which shows the function $\text{FER} = f(\text{SINR}_{\text{ref}})$ for the considered values of σ_0. We see that although the operating point $\text{FER} = 0.1$ is different for all curves, they have approximately the same shape, and the outer loop of the power control can adapt and hide most of the differences.

This means that as far as TCP-relevant properties are concerned, the outer loop power control seems to be robust to uncertainties in the radio model, in particular, to uncertainties about the performance of the inner loop power control.

The most important characteristic of the channel, at this level of modeling, is the slope of $f(r)$ around the operating point, and the step size Δ in relation to this slope. Let r^* be the operating point, with slope $f'(r^*)$. Introduce the channel characteristic $\Delta_{\text{FER}} = \Delta f'(r^*)$, which is the FER-change corresponding to one Δ-step. It seems like much of the essential system behavior is captured by this one-dimensional characteristic. Our power control, with constant Δ, can thus be expected to be less robust with respect to uncertainties that affect the slope $f'(r^*)$.

6. CONCLUSIONS

We have investigated TCP performance over WCDMA links, for varying radio model parameters, in particular performance parameters of the inner loop power control. Our result indicates that the outer loop of the power control is robust, in the sense that uncertainties in the radio model does not affect TCP/IP-layer properties.

REFERENCES

3GP (n.d.). 3rd generation partnership project; technical specification group radio access network; physical layer procedures (FDD). TS 25.214 V6.1.0.

Abouzeid, A. A., S. Roy and M. Azizoglu (2000). Stochastic modeling of TCP over lossy links. In: *INFOCOM*. Vol. 3. pp. 1724–1733.

Bai, Y., A. Rudrapatna and A. T. Ogielski (2000). Performance of TCP/IP over IS-2000 based CDMA radio links. In: *Proc. of IEEE 52th VTC'2000-Fall*. IEEE.

Barman, D., I. Matta, E. Altman and R. El Azouzi (2004). TCP optimization through FEQ, ARQ and transmission power tradeoff. In: *2nd international Conference on Wired/Wireless Internet Communications WWIC*.

Canton, A. and T. Chahed (2001). End-to-end reliability in UMTS: TCP over ARQ. In: *Globecom 2001*.

Chahed, T., A.F. Canton and S.E. Elayoubi (2003). End-to-end TCP performance in W-CDMA/UMTS. In: *IEEE ICC*.

Fischione, C., F. Graziosi and F. Santucci (2002). Outage performance of power controlled DS-CDMA wireless systems with heterogeneous traffic sources. *Kluwer Wir. Pers. Comm.* **24**, 171–187.

Fischione, C., F. Graziosi and F. Santucci (2004). Approximation of a sum of OnOff LogNormal processes with wireless applications. In: *Proc. of IEEE ICC 2004*.

Fischione, C., F. Graziosi, F. Santucci, N. Möller, K. H. Johansson and H. Hjalmarsson (2005). Modelling of TCP over wireless under power control, MAI, link level error recovery and mixed traffic sources. In: *ICC 2005*. submitted.

Gunnarsson, Fredrik and Fredrik Gustafsson (2002). Power control in wireless communications networks from a control theory perspective. IFAC World Congress, Barcelona.

Hossain, E., D. I. Kim and V. K. Bhargava (2004). Analysis of TCP performance under joint rate and power adaptation in cellular WCDMA networks. *IEEE Trans. on Wireless Comm.*

Jacobson, V. (1988). Congestion avoidance and control. In: *Proc. of SIGCOMM*. Vol. 18.4. pp. 314–329.

Khan, F., S. Kumar, K. Medepalli and S. Nanda (2000). TCP performance over CDMA2000 RLP. In: *Proc. IEEE 51st VTC'2000-Spring*. pp. 41–45.

Liu, Benyan, Dennis L. Goeckel and Don Towsley (2002). TCP-cognizant adaptive forward error correction in wireless networks. In: *Proc. of IEEE Globecom*.

Ludwig, R. and R. H. Katz (2000). The Eifel algorithm: Making TCP robust against spurious retransmissions. *ACM Computer Communication Review*.

Lundquist, Henrik and Gunnar Karlsson (2002). TCP with forward error correction. Technical report. KTH, Royal Institute of Technology, Sweden. `http://www.it.kth.se/~hen/lundqvist_tcpfec.pdf`.

Mascolo, S., C. Casetti, M. Gerla, M. Y. Sanadidi and R. Wang (2001). TCP Westwood: bandwidth estimation for enhanced transport over wireless links. In: *MobiCom*. Rome, Italy.

Möller, N. and K. H. Johansson (2003). Influence of power control and link-level retransmissions on wireless TCP. In: *Quality of Future Internet Services*. Vol. 2811 of *Lecture Notes in Computer Science*. Springer-Verlag.

Pan, J., J.W. Mark and X. Shen (2002). TCP performance and behaviors with local retransmissions. *Journal of Super Computing (special issue on Wireless Internet: Protocols and Applications)* **23**(3), 225–244.

Rossi, M., R. Vicenzi and M. Zorzi (2004). Accurate analysis of TCP on channels with memory and finite round-trip delay. *IEEE Trans. on Wireless Comm.* **3**(2), 627–640.

Sampath, Ashwin, P. Sarath Kumar and Jack M. Holtzman (1997). On setting reverse link target SIR in a CDMA system. In: *Proc. IEEE Vehicular technology conference*.

Santucci, F., G. Durastante, F. Graziosi and C. Fischione (2003). Power allocation and control in multimedia CDMA wireless systems. *Kluwer Telecommunication Systems* **23**, 69–94.

Xylomenos, G., C. Polyzos, P. Mähönen and M. Saaranen (2001). TCP performance issues over wireless links. *IEEE Communications Magazine* pp. 52–58.