Performance of TCP on a Wideband CDMA Air Interface

Michele Zorzi, Michele Rossi, and Gianluca Mazzini

University of Ferrara, 44100 Ferrara, Italy
{mzorzi, mrossi, gmazzini}@ing.unife.it
http://www.ing.unife.it

Abstract. We present a study on the performance of TCP, in terms of both throughput and energy consumption by considering a Wideband CDMA radio interface. The results show that the relationship between throughput and average error rate is largely independent on the network load, making it possible to introduce a universal throughput curve, which gives throughput predictions for each value of the user error probability. Furthermore, the possibility to select an optimal power control threshold to maximize the tradeoff between throughput and energy, is discussed.

1 Introduction

Based on the evolutionary trend of the telecommunications market, a significant step has been taken by the wireless industry, by developing a new generation of systems whose capabilities are intended to considerably exceed the very limited data rates and packet handling mechanisms provided by current second generation systems.

The goal of third generation systems is to provide multimedia services to the wireless terminals and to offer transport capabilities based on an IP backbone and on the Internet paradigm calls for the extension of widespread Internet protocols to the wireless domain.

In particular, extension of the Transmission Control Protocol (TCP) [4] has received considerable attention in recent years, and many studies have been published in the open literature, which address the possible performance problems that TCP has when operated over a connection comprising wireless links, and propose solutions to those problems (see for example [5–10]).

When TCP is run in a wireless environment, two major considerations must be made regarding its performance characterization: wireless links exhibit much poorer performance than their wireline counterparts and the effects of this behavior are erroneously interpreted by TCP, which reacts to network congestion every time it detects packet loss, even though the loss itself may have occurred for other reasons.

Furthermore, it has been shown that the statistical behavior of packet errors has a significant effect on the overall throughput performance of TCP, and that different higher-order statistical properties of the packet error process may lead to vastly different performance even for the same *average* packet error rates [11]. It is therefore important to be able to accurately characterize the actual error process as arising in the specific environment under study, as simplistic error models may just not work.

Another critical factor to be considered when wireless devices are used is the scarce amount of energy available, which leads to issues such as battery life and battery recharge, and which affects the capabilities of the terminal as well as its size, weight and cost. Since dramatic improvements on the front of battery technology do not seem likely, an approach which has gained popularity in the past few years consists in using the available energy in the best possible way, trying to avoid wasting power and to tune protocols and their parameters in such a way as to optimize the energy use [12].

It should be noted that this way of thinking may lead to completely different design objectives or to scenarios in which the performance metrics which have traditionally been used when evaluating communications schemes become less important than energy-related metrics. This energy-centric approach therefore gives a different spin to performance evaluation and protocol design, and calls for new results to shed some light on the energy performance of protocols.

Studies on the energy efficiency of TCP have been very limited so far [13–16]. In addition, they do not address specifically the Wideband CDMA environment typical of third generation systems, and therefore do not necessarily provide the correct insight for our scenario.

The purpose of this paper is to provide a detailed study of the performance of TCP, in terms of both throughput and energy, when a Wideband CDMA radio interface is used. In particular, the parameters of the UMTS physical layer will be used in the study. As a first step of the study, we present results in the absence of link-layer retransmissions. This is done in order to more clearly understand the interactions between TCP's energy behavior and the details of the radio interface. Results on the throughput performance of TCP in the presence of link-layer retransmissions can be found, e.g., in [9, 10, 17], where on the other hand energy performance is not studied.

2 System model

2.1 Simulation at the system level

In order to carry out the proposed research, a basic simulation tool has been developed. It simulates the operation of a multicellular wideband CDMA environment, where user signals are subject to propagation effects and transmitted powers are adjusted according to the power control algorithm detailed in the specifications.

We consider a hexagonal cell layout with 25 cells total. This structure is wrapped onto itself to avoid border effects. Each simulation is a snapshot in time, so that no explicit mobility is considered while running the simulation.

The fact that users may be mobile is taken into account in the specification of the Doppler frequency, which characterizes the speed of variation of the

Rayleigh fading process. For the same reason, long-term propagation effects, namely path loss and shadowing, are kept constant throughout each simulation. Path loss relates the average received power to the distance between the transmitter and the receiver, according to the general inverse-power propagation law $P_r(r) = Ar^{-\beta}$, where in our results we chose $\beta = 3.5$. Shadowing is modeled by a multiplicative log-normal random variable with dB-spread $\sigma = 4$ dB, i.e., a random variable whose value expressed in dB has Gaussian distribution with zero mean and standard deviation equal to 4 [18].

At the beginning of each simulation, all user locations are randomly drawn with uniform distribution within the service area, and the radio path gains towards all base stations are computed for each. Users are then assigned to the base station which provides the best received power. Such assignment does not change throughout the simulation. Also, note that each user is assigned to a single base station, i.e., soft handover is not explicitly considered here. Extension of the program to include soft handover is currently in progress, although qualitatively similar results are expected.

During a simulation run, the fading process for each user is dynamically changed, according to the simulator proposed by Jakes [18] and to the selected value of the Doppler frequency. Note that the instantaneous value of the Rayleigh fading does not affect the base station assignment. In order to take into account the wideband character of the transmitted signals, a frequency-selective fading model is used, with five rays whose relative strengths are taken from [19]. Maximal ratio combining through RAKE receiver is assumed at the base station. Only the uplink is considered in this paper, although similar results have been obtained for the downlink as well.

Each connection runs its own power control algorithm as detailed in the specifications. The resulting transmitter power levels of all user, along with the instantaneous propagation conditions, determine the received signal level at each receiver, and therefore the level of interference suffered by each signal. We assume here perfect knowledge of the Signal-to-Interference Ratio (SIR) which is used to make the decision about whether the transmitted power should be increased or decreased by the amount Δ. A finite dynamic range is assumed for the power control algorithm, so that under no circumstances can the transmitted power be above a maximum or below a minimum value. The delay incurred in this update is assumed to be one time unit (given by the power control frequency of update), i.e., the transmitted power is updated according to the SIR resulting from the previous update. The effect of late updates on the overall performance has been studied in [25].

Table 1 summarizes the various parameters used in the simulation.

2.2 Block error probability

The output of the system level simulations is a log of the values of the SIR, transmitted power and fading attenuation for all users, which allows us to gain some understanding on the time evolution of these quantities, as well as on the

PARAMETER	VALUE
Cell Side	200 m
β (path loss model)	3.5
A (path loss model)	-30 dB
Max. TX Power	-16 dBW
Power Range	80 dB
σ (shadowing)	4 dB
Time Unit	0.667 ms
number of Oscillators (Jakes)	8
n_rays (Selective Channel)	5
Chip Rate	3.84 Mcps
Data Rate	240 kbps
SF (Spreading Factor)	16
Δ (Power Control Step)	0.5 dB
Noise	-132 dBW

Table 1. System-level simulation parameters

behavior of the network at the system level. A post-processing package translates the SIR traces into sequences of block error probabilities (BEP).

This is done while taking into account how the radio frames are deinterleaved and decoded. The interleaving schemes have been taken from the specifications, and a convolutional code with rate 1/2 and constraint length 8 with Viterbi decoding has been considered. An analytical approximation has been used to relate a string of SIR values (one per time unit) to the probability that the corresponding block (transmitted within one or more radio frames) is in error.

The resulting trace of the block error probabilities can then be used in the simulation of higher-layer protocols, as is done in this paper, or to perform some statistical analysis of block errors. The latter approach is explored in [26], where the burstiness of the error process is looked at in some detail.

2.3 TCP simulation

For each simulation run (which corresponds to a given set of parameters), the SIR traces of all users are produced, which are then mapped into BEP traces as just explained. The latter are then used to randomly generate a block error sequence (BES).

The BES is generated from the BEP traces by just flipping a coin with the appropriate probability in each Time Transmission Interval (TTI), which is the time used to transmit a block. The BES obtained is then fed to the TCP simulator that uses it to specify the channel status in each TTI. By doing so, the SIR traces generated for all users in a simulation run are used by the TCP simulator to compute the throughput.

The average throughput is defined as the fraction of the channel transmission rate (considered at the IP level output) which provides correct information bits at the receiver, i.e., not counting erroneous transmissions, spurious retransmissions, idle time and overhead.

In addition to the throughput value, the TCP simulator computes other parameters of interest, i.e.: the average block error probability, P_e (obtained for every user simply by averaging all the values of its BEP sequence); the average energy spent in transmitting data (obtained from the transmit power traces).

These traces have been generated assuming continuous channel transmission and by updating the transmitted power according to the power control algorithm. However, when TCP is considered, transmission is bursty, due to the window adaptation mechanism implemented by the TCP algorithms, i.e., idle times occur when the window is full or the system is waiting for a timeout.

To account for idle times, we have considered the *actual* transmitted power, equal to the one obtained from power traces when TCP transmits, and to zero otherwise. The average transmitted power is then computed by summing the actual transmitted power throughout the simulation (all slots) and dividing it by the total simulation time (in number of slots). Moreover, to obtain the average consumed energy per correctly received bit, we simply divide the average transmit power by the correct information bit delivery rate.

In the following, we report the throughput and average consumed energy expressions:

$$S = \frac{CIB}{CBR}$$

$$ACE = \frac{ATP}{CIB} = \frac{ATP}{S \cdot CBR}$$

where S=average throughput, CIB=correct information bits per second, CBR= channel bit rate at the IP level output, ACE=average consumed energy per bit, and ATP=average transmitted power. We remark that, with our definitions, ACE is the average consumed energy per *correctly received bit*, i.e., the energy cost of delivering a single bit to the destination. Notice that this quantity is equal to the inverse of the *energy efficiency* of the protocol as defined in [27].

The TCP simulator implements fragmentation of TCP segments, window adaptation and error recovery. It simulates a simple unidirectional ftp session, where the direct link packet generation is assumed continuous as in a long ftp transfer. The TCP algorithm considered is *New Reno* [21]. Data flow is unidirectional, i.e., data packets flow only from sender to receiver, while ACKs flow in the reverse direction. Receiver generates non-delayed ACKs, i.e., one ACK is sent for each packet received. The TCP/IP stack is version 4, with a total of 40 bytes (including both TCP and IP overhead) for each header (compression is not taken into account in the results presented) and MTU size of 512 bytes. We have not considered RLC and MAC levels, which are assumed here to operate in transparent mode. RLC/MAC level implementation and characterization are currently under study.

To compute the bit rate at the output of the IP level, we have to account for overhead added by the physical layer as well as possibly due to multiplexing of other channels. For the purpose of discussion, we assume the following figures: a transport block is 1050 bits, including 16 bits of overhead due to transparent RLC/MAC operation; a CRC and tail bits for code termination are added to this block, and the result is convolutionally encoded at rate 1/2. The resulting encoded block is then brought to 2400 encoded symbols by the rate matching algorithm, so that the raw physical layer symbol rate is 240 kbps. The application of a spreading factor $SF = 16$ makes it 3.84 Mcps, which is the standard channel transmission rate. At the IP level output, we then have a block of 1034 bits of data every 10 ms, thereby yielding a net bit rate of 103.4Kbps, which is the bit rate used in the TCP simulator.

The use of TCP New Reno algorithm has been motivated by its implementation of fast recovery and fast retransmit algorithms [22], as recommended in [23], especially for wireless environments. This is an optimization over previous TCP versions, and is currently at the Proposed Standard level.

3 TCP throughput performance

In the graphs presented we will indicate with N_u the number of users in the simulation, with TTI the number of radio frames over which interleaving is performed, $SIR_{th} = t + \Delta$ dB indicates that the threshold used in the power control algorithm is t, while Δ is its increment as described in Section 2. Finally, with the term f_d we refer to the Doppler frequency used in the Rayleigh fading simulator. In the following graphs, the results will be represented as average TCP throughput, S, vs. average block error probability, P_e, thereby assigning a single point in the graph to each user.

3.1 Doppler frequency effect

Figure 1 shows the TCP throughput performance for different values of the Doppler frequency. The graph is plotted by reporting throughput vs. P_e for each of the 90 users involved in the simulation; each user is identified by a marker. The case of independent errors is also reported for comparison purpose (here the markers are used only to identify the curve and are not related to the users).

Note that, for a given value of the Doppler frequency the points representing the various users of a simulation appear to lie along a fairly well-defined curve. It is worth stressing that this was not obvious a priori since different users are placed in different locations and are subject to different propagation conditions, both in terms of slow impairments (log-normal shadowing) and in terms of fast fading.

This allows us to introduce the concept of "universal throughput curve" for a given situation, in the sense that users which suffer similar values of P_e will enjoy about the same throughput. Again, this is not obvious since different users in a simulation may see different statistical behaviors of the errors, which could in principle lead to different performance even in the presence of the same average error rate [11].

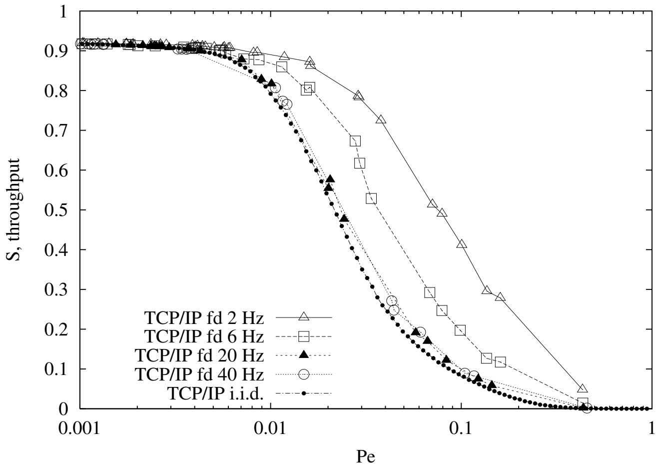

Fig. 1. TCP/IP Sensitivity to Doppler,($N_u = 90$, $TTI = 1$, $SIR_{th} = 5.5 + 0.5$dB, $f_d = 2, 6, 20, 40$Hz).

An explanation can be drawn from the results in [26], where it was found that for a given value of f_d there is a strong correlation between P_e and the error burstiness. In this situation, P_e essentially determines the full second-order characterization of the error process, which in turn almost fully specifies the value of the TCP throughput. On the other hand, for different values of f_d we observe different curves. In fact, even in the presence of the same P_e the different extent of the channel memory results in different performance.

Another interesting observation from Figure 1 is that as the Doppler frequency increases the performance degrades, i.e., slower channels correspond to better performance, as already observed in [8]. As expected, for sufficiently high values of the Doppler frequency, the behavior of the system is close to the iid case.

Finally, we note that the shape of the curves appears to be fairly regular, with a smooth transition from highest throughput values (essentially limited only by the percentage of overhead in the TCP packets, far left of the graph), to essentially zero throughput when errors are very likely (right end of the graph). This shape, which has been observed by other authors, lends itself nicely to numerical fitting, as detailed later.

3.2 Interleaving depth effect

In UMTS, besides the so-called *intra-frame* interleaving, which is always used to scramble the bits within a radio frame (10 ms) before encoding, it is also possible to use a second, *interframe*, interleaving, which mixes bits *across* frames. By doing so, the performance of the decoder is of course improved also in the presence of burst errors, but a larger interleaving delay is introduced.

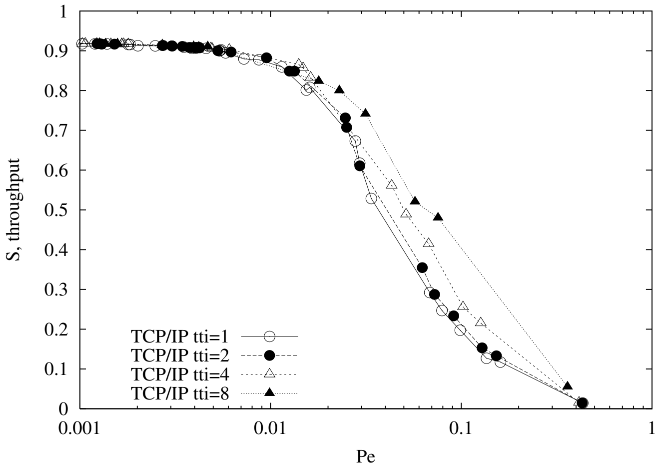

Fig. 2. Average throughput *vs* TTI,(N_u = 90, $TTI = 1, 2, 4, 8$, $SIR_{th} = 5.5 + 0.5$dB, f_d = 6Hz).

Therefore, another interesting sensitivity analysis regards the interleaving span allowed by the application. While keeping in mind the price to be paid in terms of delay (an increase of the TTI corresponds to a larger delay), we can note from Figure 2 how a deeper interleaving gives a beneficial effect. Notice also that the shape of the curves is similar to what observed in the previous case.

3.3 Network load effect

The effect of the network load is shown in Figures 3 and 4. In Figure 3, results from three simulations are shown, with 80, 90 and 100 users in the network, respectively. For each simulation, all users are assigned a point with the same abscissa (which is given by the number of users in the system) and with vertical coordinate given by their average TCP throughput.

We can see how for increasing load the presence of disadvantaged users becomes more noticeable, as expected, and the system is more and more unfair. The same results are represented in Figure 4 by reporting the throughput as a function of P_e where the curve along which the various points are aligned is relatively insensitive to the system load.

In Figure 1, it was observed that, in a given scenario, two users whose average block error probability is the same will have essentially the same throughput, regardless of the specific situation of each. What we see here is that this behavior still holds across simulations in which different levels of network load are considered (but for the same Doppler frequency). For higher network load, each user will certainly see worse performance due to the increased interference, but the relationship between average throughput and average error rate is essentially unaffected.

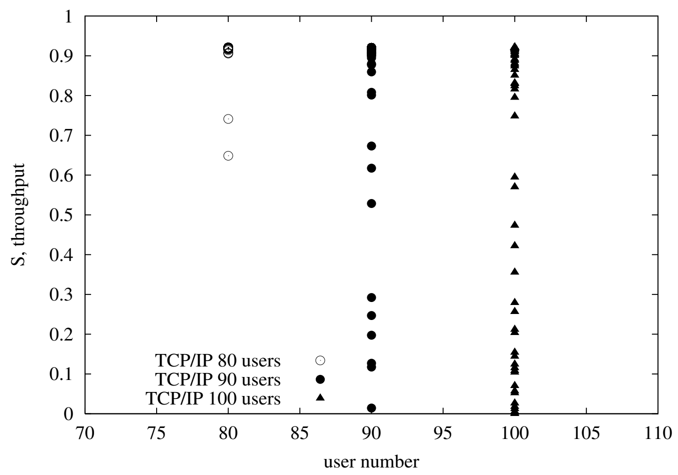

Fig. 3. TCP/IP average throughput *vs* network load, ($N_u = 80, 90, 100$, $TTI = 1$,$SIR_{th} = 5.5 + 0.5$dB, $f_d = 6$Hz).

This highlights the power of the concept of "universal curve" which can be used to study typical cases and to infer TCP throughput performance based only on easily measurable physical layer parameters. In Figure 4 a possible form for the universal curve is given for comparison with the simulator output points; more details about this fitting function will be presented below. Similar behavior has been observed for other values of the Doppler frequency.

3.4 Analytical throughput prediction

The observed shape of the throughput curves, which tend to a constant equal to one minus the percentage of overhead for $P_e \to 0$ and to zero for $P_e \to 1$, suggests a numerical fit involving the logarithm of P_e. Also, the shape of the transition is seen to depend on the value of the Doppler frequency.

In Figure 5 we show a proposal for the modeling of the TCP throughput behavior through a heuristic function $f(x)$, independent of the network load and parameterized only by the Doppler frequency, as suggested by the curves of Figure 4 as well as by other results not shown in this paper. The proposed expression for $f(x)$ is as follows:

$$f(x) = S(0) \cdot \frac{10^{\alpha \cdot \ln(\frac{1}{\tilde{x} - x_s} - 1)}}{10^{\alpha \cdot \ln(\frac{1}{\tilde{x} - x_s} - 1)} + 1}$$

with $\tilde{x} = 1 + \frac{\log_{10}(x)}{3}$, $\alpha = 1.3$, $x_s = \frac{1}{Af_d + B} - k$ and where $A = 1.39$, $B = 2.78$, $k = 0.03$, and $S(0)$ is the average throughput for $P_e = 0$ and f_d is the Doppler frequency in Hz.

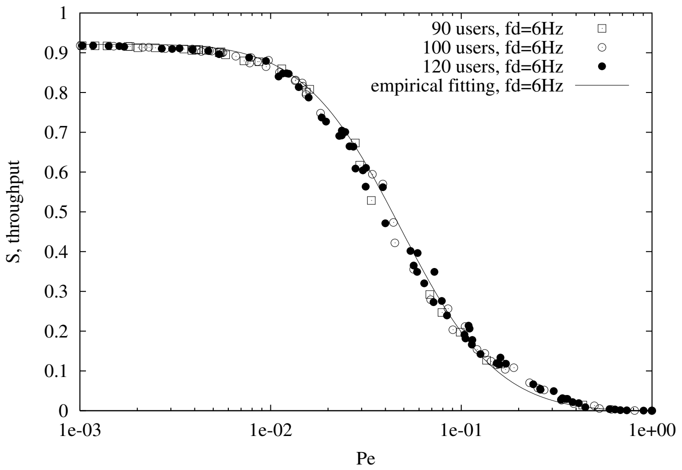

Fig. 4. Throughput curve, independence from network load with Doppler 6Hz, ($N_u = 80, 90, 100$, $TTI = 1$, $SIR_{th} = 5.5 + 0.5$dB, $f_d = 6$Hz).

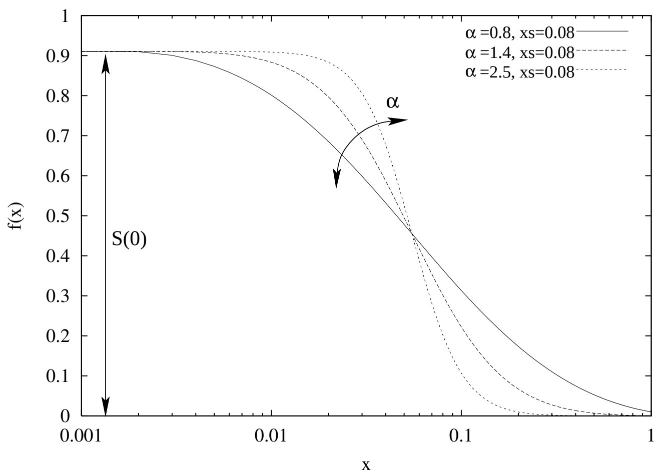

Fig. 5. Throughput heuristic function.

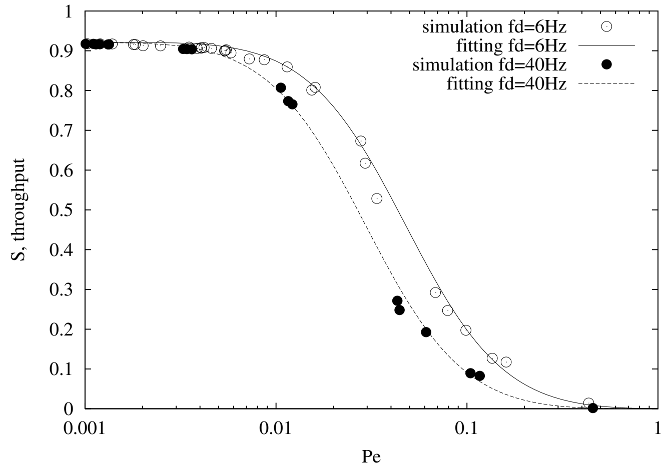

Fig. 6. Throughput heuristic function: approximation of simulator outputs, ($N_u = 90$, $TTI = 1$, $SIR_{th} = 5.5 + 0.5$, $f_d = 6, 40\text{Hz}$).

The accuracy of the proposed fit has been tested for various values of the parameters involved. Examples of these tests are given in Figures 4 and 6, in which the fitting expression is compared against the simulation results for two values of the Doppler frequency. These graphs show that the analytical expression is reasonably close to the actual points obtained by simulation.

4 TCP energy performance

All previous results were obtained for a given value of the power control threshold, SIR_{th}, which is used to drive the transmit power dynamics at each user and which directly affects the error performance. In fact, choosing a higher value of this threshold has the double effect of forcing the users to transmit more power in order to achieve a higher SIR (thereby consuming more energy) but also of causing the SIR experienced by the typical user to be higher (thereby improving the error rates and therefore the TCP throughput).

It is therefore of interest to study how varying the power control SIR threshold makes it possible to cut a tradeoff between QoS and energy consumption.

4.1 Throughput and consumed energy

Figures 7 and 8 show the trade-off between TCP throughput and consumed energy. Each curve corresponds to a given user for a given set of parameters, whereas different points on the same curve refer to different values of the power control threshold. Figure 7 shows the effect of using different threshold values in terms of S as a function of ATP. In this figure only the behavior of some

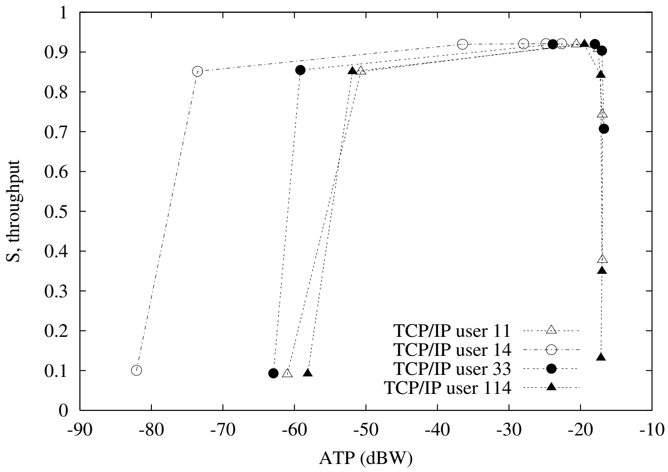

Fig. 7. Throughput *vs* average transmitted power for different threshold values, ($N_u = 120$, $TTI = 1$, $SIR_{th} = \{1.5, 2.5, 3.5, 4, 4.5, 5\} + 0.5\text{dB}$, $f_d = 6\text{Hz}$).

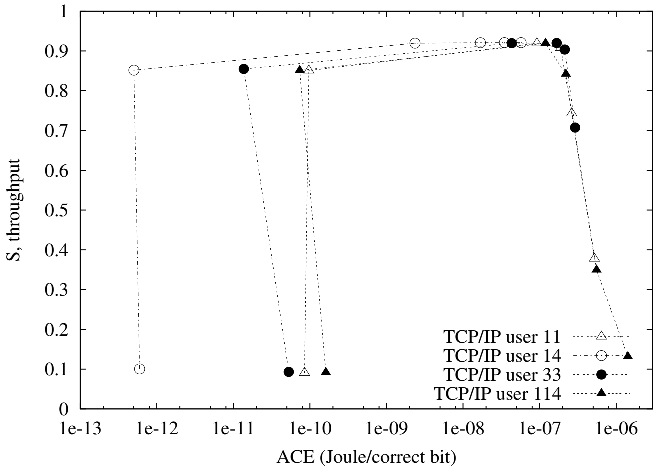

Fig. 8. Throughput *vs* average consumed energy for different threshold values, ($N_u = 120$, $TTI = 1$, $SIR_{th} = \{1.5, 2.5, 3.5, 4, 4.5, 5\} + 0.5\text{dB}$, $f_d = 6\text{Hz}$).

selected users has been reported. These users can be seen as representative of all users in the network, in the sense that they illustrate typical behaviors as they arise in the system.

In Figure 8 the same results are shown, by considering ACE instead of ATP; as expected, for low throughput the obtained curves are shifted to the right, since ACE is obtained by dividing ATP by the throughput S (except for an inessential constant scaling factor. This is intuitively explained by the fact that, for a given average consumed power, the energy per correct information bit is greater for low throughput values, i.e., when it is hard to deliver bits correctly, the cost associated to each one of them is higher. Notice that TCP already does the right thing by stopping transmission when the channel is very bad (the timeout event), whereas it tries to recover from errors whenever possible through retransmissions, which may waste some power.

In general, increasing the threshold should lead to better performance for many users, since the SIR experienced is expected to be higher; this is not necessarily true, however, since, in order to achieve a higher SIR threshold, many users will transmit more power, thereby causing more interference in the system. If the SIR objective is not achievable for all users in the system, some users will actually see degraded performance for higher values of the threshold since, although the threshold value would correspond to better performance, they cannot achieve its value.

From the obtained results we have noted different users behavior. For some users an increasing power threshold always corresponds to a greater throughput: for these users the throughput as a function of the power threshold is a monotonic curve. For others, the throughput vs. consumed power curve increases up to a breakpoint, after which an increment of the power threshold (and thereby of the transmitted power) actually leads to worse performance, due to greater interference as discussed above.

In our results, user 14 is the one showing monotonic behavior, since it experiences favorable propagation conditions, and therefore is not significantly affected by the increased interference level in the system. For users 11, 33 and 114, a different situation can be observed. In particular, user 11 is the one with the worst behavior as the power threshold increases.

In any event, from these results we can conclude that increasing the target value of the SIR in the system does not necessarily translate into improved quality, but there exists an optimal value of the threshold, beyond which some users will experience negligible throughput improvements, whereas others will even see degraded performance. In the cases studied in this paper, this optimal value is seen to be close to 3.5 dB.

4.2 Energy and threshold

Another important remark regards the numerical values shown in Figures 7 and 8. It can be clearly seen that unlike for throughput, which except for badly chosen values of the threshold exhibits relatively small variations, the range spanned by the energy performance extends over multiple orders of magnitude. This

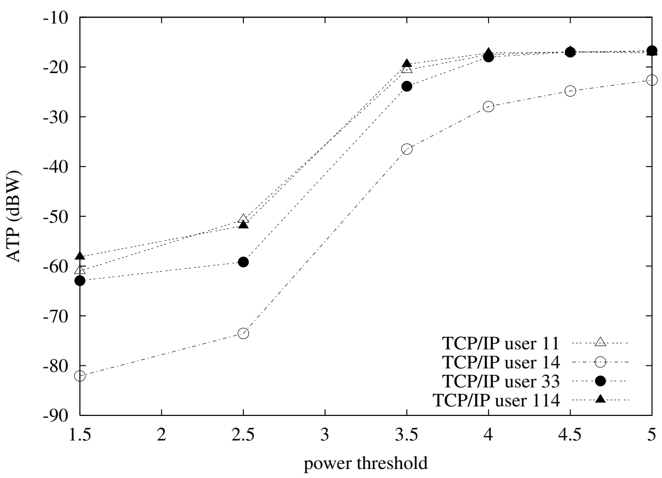

Fig. 9. Average consumed power *vs* power threshold, ($N_u = 120$, $TTI = 1$, $SIR_{th} = \{1.5, 2.5, 3.5, 4, 4.5, 5\} + 0.5\text{dB}$, $f_d = 6\text{Hz}$).

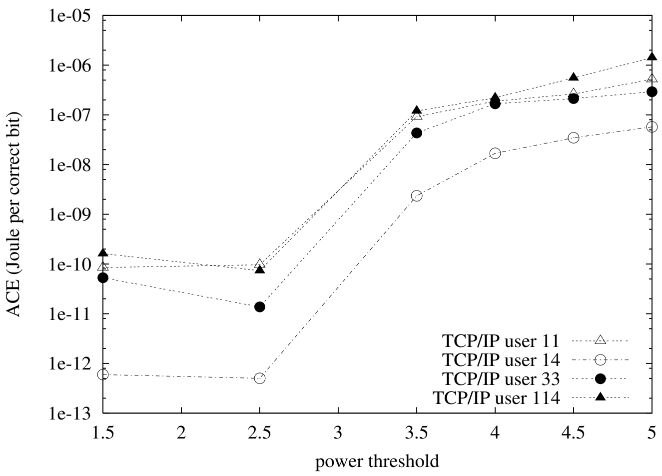

Fig. 10. Average consumed energy *vs* power threshold, ($N_u = 120$, $TTI = 1$, $SIR_{th} = \{1.5, 2.5, 3.5, 4, 4.5, 5\} + 0.5\text{dB}$, $f_d = 6\text{Hz}$).

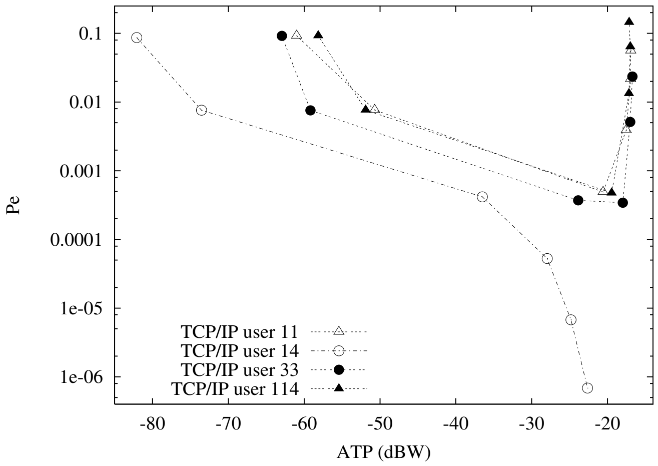

Fig. 11. P_e *vs* average transmitted power for different threshold values, ($N_u = 120$, $TTI = 1$, $SIR_{th} = \{1.5, 2.5, 3.5, 4, 4.5, 5\} + 0.5$dB, $f_d = 6$Hz).

indicates that the choice of the proper power control threshold, while certainly important for error and throughput considerations, becomes critical when energy performance is considered.

Figure 9 shows the average consumed power as a function of the power threshold. As expected, a greater threshold value always corresponds to a higher consumed power, i.e., all curves have a strictly monotonic behavior. This is due to the fact that for higher threshold values, more power is necessary in order to obtain the required SIR target.

A different behavior is observed in Figure 10 in which the consumed energy per bit is reported instead. In particular, notice that in the far left of the graph, a decrease of the threshold, although corresponding to smaller average power (see Figure 9), results in a higher energy cost per bit. This behavior corresponds to users suffering from low throughput performance, where the consumed energy per correct information bit grows, as explained before. As a last observation, we note from Figure 9 that users 11, 33 and 114, i.e., the ones that suffer from system interference as the threshold grows, show a transmitted power which is essentially constant for values of the threshold beyond 4 dB, This is due to the fact that these users, in trying to achieve the required SIR and to make up for the increased interference, have reached the maximum allowed value for the transmit power, and therefore their power can not be increased any more.

4.3 Error probability and consumed energy

A similar QoS-energy trade-off relates to the error rate performance instead of the TCP throughput. In Figure 11 the block error probability P_e has been reported against the average transmitted power ATP by using different values of

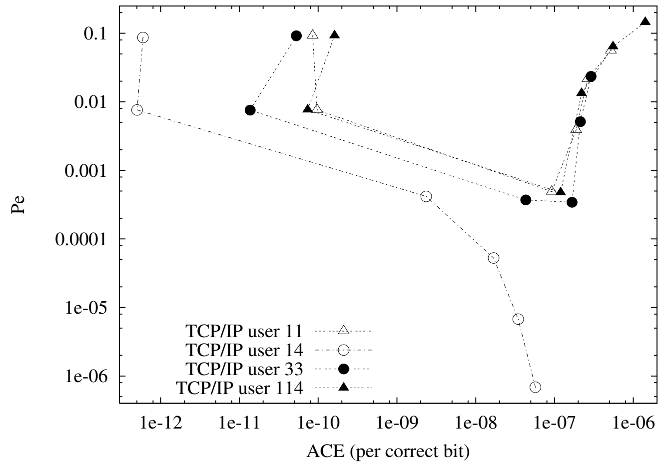

Fig. 12. P_e *vs* average consumed energy for different threshold values, ($N_u = 120$, $TTI = 1$, $SIR_{th} = \{1.5, 2.5, 3.5, 4, 4.5, 5\} + 0.5$dB, $f_d = 6$Hz).

the threshold. From the graph, we note that P_e decreases as the threshold grows until a minimum is reached. After this point P_e starts to grow, again due to the increased interference in the system. The points in which P_e has a minimum are the same on which the throughput of the system is maximized (see Figure 7).

As before, user 14 is the only user considered for which P_e never grows, i.e., increasing the threshold always leads to better performance. In Figure 12 ACE is reported instead of ATP and, as in the previous cases, some points are shifted to the right due to poor throughput performance.

From the above results, we may conclude that the selection of the power control SIR threshold is critical in cutting the right trade-offs between QoS and energy performance. In particular, we observe that the potential for energy gains is very significant compared to similar effects on throughput, being measured over multiple orders of magnitude. Therefore, it seems that more attention should be given to energy consumption issues at the Radio Resource Management level, which is responsible for the power control parameter selection.

5 Conclusions

In this paper, some results on the behavior of TCP over a Wideband CDMA air interface have been reported. In particular, TCP throughput curves have been obtained, and their dependence on various parameters, such as the number of users, the interleaving depth, and the Doppler frequency, has been investigated.

From this study, we have found that the relationship between average TCP throughput and average block error rate is largely independent of the number of users in the system. For this reason, it is possible to empirically characterize such a curve with a matching function that only depends on the Doppler frequency. A study of the energy consumption has also been performed, showing that a trade-off between the throughput and the power control threshold exists. It is therefore possible to trade-off QoS of the data transfer for increased energy efficiency. For many users, it has been shown that an optimal value of the threshold exists, potentially leading to very significant energy savings in return for very small throughput degradation. In the considered case, values of the power control threshold close to 3.5-4 dB cut the best tradeoff.

In order to focus exclusively on the interactions of TCP with the WCDMA radio technology and as a preliminary step in characterizing TCP's energy performance in this scenario, no link-layer retransmissions have been considered in this study. Future work includes extension of the study to the presence of a radio link layer which improves the wireless link performance through block retransmission. Similar performance studies for TCP versions other than New-Reno also seem worth pursuing.

References

1. G. Patel, S. Dennett. The 3GPP and 3GPP2 movements toward an all-IP mobile network. *IEEE Personal Communications*, vol. 7, pp. 62-4, Aug. 2000.
2. B. Sarikaya. Packet mode in wireless networks: overview of transition to third generation. *IEEE Communications Magazine*, vol. 38, pp. 164-72, Sept. 2000.
3. Third Generation Pasrtnership Project website: www.3gpp.org
4. W. R. Stevens. *TCP/IP Illustrated, Volume 1.* Addison Wesley, 1994.
5. H. Balakrishnan, V. N. Padmanabhan, S. Seshan, and R. H. Katz. A comparison of mechanisms for improving TCP performance over wireless links. *ACM/IEEE Trans. on Networking*, Dec. 1997.
6. R. Caceres, and L. Iftode. Improving the performance of reliable transport protocols in mobile computing environments. *IEEE Jl. Sel. Areas Commun.,* vol. 13, no. 5, pp. 850-857, June 1995.
7. A. Kumar. Comparative performance analysis of versions of TCP in a local network with a lossy link. *ACM/IEEE Trans. on Networking*, Aug. 1998.
8. M. Zorzi, A. Chockalingam, R.R. Rao. Throughput Analysis of TCP on Channels with Memory. *IEEE J. Selected Areas Comm.*, vol. 18, pp. 1289–1300, Jul. 2000.
9. A. Chockalingam, Gang Bao. Performance of TCP/RLP protocol stack on correlated fading DS-CDMA wireless links. *IEEE Trans. on Vehicular Technology*, vol. 49, pp. 28-33, Jan. 2000.
10. N.M. Chaskar, T.V. Lakshman, U. Madhow. TCP over wireless with link level error control: analysis and design methodology. *IEEE/ACM Trans. on Networking*, vol. 7, pp. 605-15, Oct. 1999.
11. M. Zorzi, R.R. Rao. Perspectives on the Impact of Error Statistics on Protocols for Wireless Networks. *IEEE Personal Communications*, vol. 6, Oct. 1999.
12. *IEEE Personal Communications Magazine*. special issue on "Energy Management in Personal Communications and Mobile Computing," vol. 5, June 1998.
13. M. Zorzi, R.R. Rao. Is TCP energy efficient?. in *Proc. IEEE MoMuC*, Nov. 1999.
14. V. Tsaoussidis, H. Badr, X. Ge, K. Pentikousis. Energy/throughput tradeoffs of TCP error control strategies. In *Proceedings ISCC 2000* pp. 106-12, Jul. 2000.

15. V. Tsaoussidis, A. Llahanns, H. Badr. Wave & wait protocol (WWP): high throughput and low energy for mobile IP-devices. In *Proceedings IEEE ICON 2000*, pp. 469-73, Sept. 2000.
16. M. Zorzi, R.R. Rao. Energy Efficiency of TCP. To appear in *Mobile Networks and Applications*, 2001.
17. F. Khan, S. Kumar, K. Medepalli, S. Nanda. TCP performance over cdma2000 RLP. In *Proc. VTC2000-Spring*, pp. 41-5, May 2000.
18. W. C. Jakes, Jr. *Microwave mobile communications.* New York: John Wiley & Sons, 1974.
19. ETSI SMG2. Evaluation report for ETSI UMTS Terrestrial Radio Access (UTRA) ITU-R RTT candidate. Sep. 1998 (http://www.itu.int/imt/2-radio-dev/reports/etsi/eva_utra.pdf)
20. Harri Holma, Anttiq Toskala. *WCDMA for UMTS: Radio Access For Third Generation Mobile Communications.* New York: John Wiley & Sons, 2000.
21. S. Floyd, T. Henderson. The New Reno Modification to TCP's Fast Recovery Algorithm. Request For Comment 2582, April 1999.
22. M. Allman, V. Paxson, W. Stevens. TCP congestion control. Request For Comment 2581, April 1999.
23. S. Dawkins, G. Montenegro, M. Kojo, V. Magret, N. Vaidya. End-to-end Performance Implications of Links with Errors. Internet Draft, November 17, 2000.
24. Anna Calveras, Josep Paradells Aspas. TCP/IP over wireless links: Performance Evaluation. VTC'98. 48th IEEE. Volume: 3, 1998, pages: 1755-1759.
25. A. Giovanardi, G. Mazzini, V. Tralli, M. Zorzi. Some results on power control in Wideband CDMA cellular networks. In *Proc. WCNC2000*, Chicago (IL), Sep. 2000.
26. M. Zorzi, G. Mazzini, V. Tralli, A. Giovanardi. Some results on the error statistics over a WCDMA air interface. In *Proc. MMT2000*, Florida (USA), Dec. 2000.
27. M. Zorzi, R.R. Rao. Energy constrained error control for wireless channels. *IEEE Personal Communications Magazine*, vol. 4, pp. 27–33, Dec. 1997.