Performance of TCP/RLP Protocol Stack on Correlated Fading DS-CDMA Wireless Links

A. Chockalingam, *Senior Member, IEEE*, and Gang Bao, *Member, IEEE*

***Abstract*—In this paper, we present the throughput performance of transport control protocol/radio link protocol (TCP/RLP) stack on correlated fading direct-sequence code-division multiple-access (DS-CDMA) wireless links. It is shown that because of significant burstiness in RLP frame errors in highly correlated Rayleigh fading, longer persistence at the RLP layer to recover lost RLP frames (more than the IS-99 specified three retransmission attempts at the RLP layer) is beneficial at low-link fading margins.**

***Index Terms*—CDMA, fading, RLP, TCP.**

I. INTRODUCTION

DUE TO rapid advances in the area of wireless communications and internet, provision of data services for applications like e-mail, web browsing, telnet, etc., over wireless is gaining importance. transport control protocol (TCP) is a reliable end-to-end transport protocol in the internet protocol (IP) suite [1]. TCP is tuned to perform well in wireline networks where the channel error rates are very low and congestion is the primary cause of packet loss [2]. However, when used over wireless channels which are characterized by high frame error rates (FER's), the performance of TCP is severely affected [3], [4]–[6]. In order to reduce the FER seen by the TCP layer on wireless links, a radio link protocol (RLP) is generally introduced at the link layer, above the physical layer.

Typically, the RLP uses an ARQ error recovery mechanism to reduce the FER. For example, in direct-sequence code-division multiple-access (DS-CDMA) cellular systems, a RLP with a negative acknowledgment (NAK)-based selective repeat scheme has been standardized (IS-99 standard) [7]. In IS-99 data services standard, whose underlying physical layer is defined by IS-95 standard using DS-CDMA [8], the RLP performs a partial link layer recovery through a limited number of RLP frame retransmissions in case of frame error. In the event of RLP failure due to excessive frame errors, control is passed on to the TCP layer which is ultimately responsible for providing complete end-to-end recovery. The ARQ mechanisms at RLP and TCP layers thus interact with each other and affect the overall system performance depending on the channel conditions. Performance of transport control protocol/radio link protocol (TCP/RLP) over DS-CDMA wireless links has been studied in [9]. However, one of the assumptions in the performance evaluation in [9] is that the frame error distribution is independent and identically distributed (i.i.d.). This assumption may not be valid when the link suffers from *correlated* multipath fading because of which frame errors could be bursty. The length of the frame error bursts depends on the link fading margin and the Doppler bandwidth of the fading process [10]–[12]. The smaller the Doppler bandwidth (equivalently, slower the mobile speed), the longer the error bursts will be.

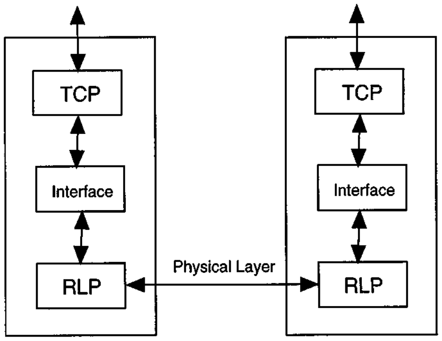

Fig. 1. Typical TCP/RLP protocol stack.

In this paper, we are concerned with the performance of TCP/RLP in the presence of bursty frame errors caused by correlated multipath fading. Specifically, we consider a situation where a bulk data transfer is being performed over a TCP connection from the base station interworking function (IWF) to a mobile terminal over a DS-CDMA link carrying 20-ms RLP frames (192 bits per frame at 9600-bps transmission rate). We provide a characterization of TCP/RLP throughput as a function of the Doppler bandwidth and the average FER. We do this by modeling the frame success/failure process over multipath fading channels as a first-order Markov process [11], [12]. We illustrate interesting TCP/RLP performance behavior under various channel fading scenarios.

II. TCP/RLP PROTOCOL STACK

A typical TCP/RLP protocol stack in a wireless environment is shown in Fig. 1. The full protocol stack in IS-99 [7], however, includes network (IP) and link (PPP) layers between the TCP and RLP layers. We do not show the IP and PPP layers here because these layers generate a relatively fixed amount of overhead to the overall performance. In other words, the performance of the stack is mostly determined by the interaction between the TCP and RLP layers under varying channel conditions. The TCP data unit, called a *segment*, is typically 600 bytes in length. As per IS-99, the maximum segment size (MSS) at the TCP layer is no smaller than 536 bytes (default value), the

---

Manuscript received January 23, 1998; revised November 23, 1998. This paper was presented at the IEEE Vehicular Technology Conference, Ottawa, Canada, May 1998.

The authors are with Qualcomm, Inc., San Diego, CA 92121-2779 USA (e-mail: achockal@qualcomm.com).

Publisher Item Identifier S 0018-9545(00)00638-1.

transmitted TCP segment size is no larger than 2047 bytes, and the advertised window size is no smaller than 2*MSS and no larger than 4*MSS. Successful transmission of each TCP segment is ensured through an acknowledgment based scheme [1]. The RLP data unit, called a *frame*, on the other hand, is about 20 bytes in length (equivalent to a fixed frame duration of 20 ms at a transmission rate of 9600 bps). Hence, a typical TCP segment transmission requires about 30 RLP frame transmissions. When transferring data, RLP is a pure NAK-based selective repeat protocol [7]. That is, the receiver does not acknowledge correct RLP data frames; it only requests the retransmission of RLP data frames that were not received correctly. In case of a frame error, RLP performs a partial link recovery through a limited number of frame retransmissions. Since the number of allowed retransmissions of an erroneous RLP frame is finite, RLP can not completely eliminate all detectable errors (undetectable errors, which in general can be bounded to a very small value by a properly chosen CRC, are not considered here). Unrecovered errors at the RLP layer will be recovered by the TCP layer.

When RLP at the receiving end finds a frame in error (or as missing), it sends back a NAK requesting retransmission of the lost frame. A NAK retransmission timer is set for the lost frame. A guard interval of five frames (i.e., 100 ms) is added to the retransmission timeout in order to account for the buffering delays and the segmentation of retransmitted frames [7]. When the retransmission timer expires for the first time, RLP resets the timer and sends back NAK twice. Each NAK received at the transmission end triggers a retransmission of the requested frame. When the timer expires for the second time, RLP resets the timer and sends back three NAK's. This process continues until the number of timer expirations reaches a certain limit n ($n = 3$ by default in IS-99). RLP will abort the attempt after n unsuccessful retransmissions and pass control to the TCP layer.

During the data transfer phase (i.e., excluding the link setup and tear-down phases), RLP maintains a 8-b sending sequence number count $V(S)$, and two sequence numbers $V(R)$ and $V(N)$ for receiving. All operations on these RLP frame sequence numbers are carried out in unsigned modulo 256 arithmetic. $V(S)$ is incremented whenever a new RLP data frame of nonzero bytes is sent out. That is, $V(S)$ is the sequence number of the next frame to be sent. $V(R)$ contains the sequence number of the next new frame expected to be received. $V(N)$ is the next frame needed for sequential delivery. In other words, $V(N)$ is the oldest sequence number of the missing frames. By denoting the sequence number of a newly received frame by i, the RLP transmission procedure can briefly be described by the following rules [7].

- If $i < V(N)$ or if the frame is already stored in the resequencing buffer, discard the frame.
- If $i = V(N)$, update $V(N)$ to the next oldest missing frame sequence number. Pass received frames up to $V(N) - 1$ to the upper layer.
- If $V(N) < i < V(R)$, store frame i in resequencing buffer if it is missing.
- If $i = V(R) = V(N)$, pass all received frames up to $V(R)$ to the upper layer.
- If $i = V(R) \neq V(N)$ or $i > V(R)$, increment $V(R)$ and store frame i into resequencing buffer.
- For all cases, send NAK's of missing frames if their retransmission timers are not yet set or expired.

III. CORRELATED FADING CHANNEL MODEL

We consider a DS-CDMA forward link (base station-to-mobile link) where the slowly varying shadow and distance losses are perfectly compensated, and the rapid variations due to multipath fading remain uncompensated. The multipath fading on a mobile radio channel is considered to follow a Rayleigh distribution [10]. The issue of modeling the characteristics of such a channel has been addressed in [11], where it was shown that the Rayleigh fading envelope can be well approximated by a first-order Markov process with continuous amplitude. A further step in the modeling of block error process in a correlated Rayleigh fading channel was due to [12], where it was shown that, for a broad range of parameters, the sequence of data block success and failure can itself be approximated by means of a simple two-state Markov chain whose transition probability matrix is given by

$$M_c = \begin{pmatrix} p & 1-p \\ 1-q & q \end{pmatrix} \tag{1}$$

where p and $1-q$ are the probabilities that the jth block transmission is successful, given that the $(j-1)$th block transmission was successful or unsuccessful, respectively. Given the matrix M_c, the channel properties are completely characterized. In particular, it is possible to find the steady-state distribution of the chain. The steady-state probability that a frame error occurs (P_E) is

$$P_E = \frac{1-p}{2-p-q}. \tag{2}$$

Note that $(1-q)^{-1}$ represents the average length of a burst of frame errors, which is described by a geometric random variable. Also, for a Rayleigh fading channel with a fading margin F, the average FER can be found as [12]

$$P_E = 1 - e^{-1/F} \tag{3}$$

and the Markov parameter q can be derived as [12]

$$q = 1 - \frac{Q(\theta, \rho\theta) - Q(\rho\theta, \theta)}{e^{1/F} - 1} \tag{4}$$

where

$$\theta = \sqrt{\frac{2/F}{1-\rho^2}}. \tag{5}$$

In the above, $\rho = J_0(2\pi f_d T)$ is the Gaussian correlation coefficient of two successive samples of the complex amplitude of a fading channel with Doppler frequency f_d, taken T s apart. The parameter $f_d T$ is the normalized Doppler bandwidth which describes the correlation in the fading process. $J_0(\cdot)$ is the Bessel function of the first kind and zero order. $Q(\cdot, \cdot)$ is the Marcum Q function, given by

$$Q(x, y) = \int_y^{\infty} e^{-\frac{(x^2+w^2)}{2}} I_0(xw) w dw \tag{6}$$

where I_0 is the modified Bessel function of the first kind and of zeroth order. From P_E and q in (3) and (4), the Markov parameter p can be obtained using (2). In this paper, we apply this

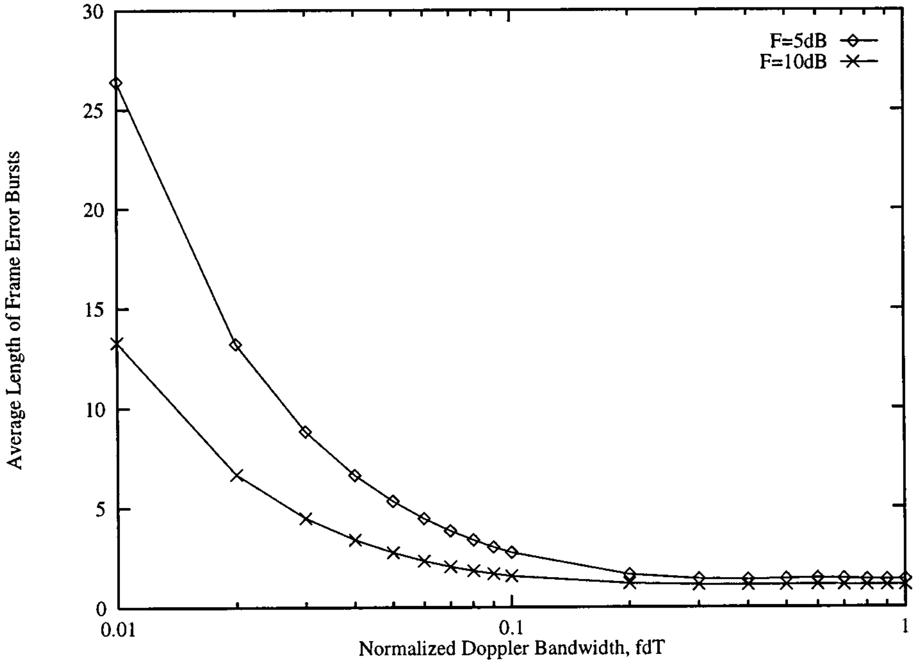

Fig. 2. Average length of frame error bursts $(1-q)^{-1}$ versus normalized Doppler bandwidth f_dT at different link fading margins F.

Markov model at the RLP frame level (i.e., $T = 20$ ms). By choosing different P_E and f_dT values, we can establish fading channel models with different degree of correlation in the fading process. Fig. 2 shows the average length of frame error bursts, $(1-q)^{-1}$, as a function of the normalized Doppler bandwidth f_dT and the link fading margin F. When f_dT is small, the fading process is very correlated (long bursts of frame errors); on the other hand, for larger values of f_dT, successive samples of the channel are almost independent (short bursts of frame errors).

IV. RESULTS AND DISCUSSION

We consider a scenario where a large data file is transferred from the base station host to a mobile terminal over a DS-CDMA wireless link. The performance measures of interest here are the end-to-end throughput at the RLP and TCP layers. RLP throughput refers to the amount of successfully transferred RLP frames in steady state, excluding the overhead bits in the RLP frames. For example, a 20-ms full-rate RLP frame in IS-99 consists of 168 information bits and 24 overhead bits. Therefore, at 9600-bps transmission rate, the achievable RLP throughput is just 8000 bps, even if the channel is error free (i.e., no RLP frame errors). RLP frame losses will degrade the throughput further (e.g., at 1% FER, the resulting RLP throughput would be 7920 bps). TCP throughput, on the other hand, refers to the amount of successful information bits transferred at the TCP layer in steady state. Not all correctly received RLP frames contribute to the throughput at the TCP layer because even if one RLP frame is lost in a TCP segment, the entire TCP segment is lost. Consequently, the throughput achieved at the TCP layer is typically lower than the RLP throughput. Note that even when there are no RLP frame errors, there will be a loss of throughput at the TCP level compared to the RLP throughput because of the TCP/IP and PPP headers. A TCP/IP header size of 3 bytes (VJ compressed) and a PPP overhead size of 4 bytes (1-byte flag, 1 header byte, 2 CRC bytes) are typical. Thus, for a 536-byte TCP segment, a minimum TCP throughput loss of 1.3% due to TCP/IP and PPP overheads is inevitable.

The TCP throughput performance depends not only on the average FER at the RLP layer, but also on the amount of burstiness in the RLP frame errors. In this section, we evaluate the TCP/RLP throughput performance as a function of both the average FER as well as the average length of frame error bursts. In a typical DS-CDMA link, an average FER of 1% is the typical operating point. From (3), a 1% average FER (i.e., $P_E = 0.01$) corresponds to a link fading margin (F) of about 20 dB. The range of average FER values we consider here is from 1% to 30%. Note that 10% and 30% average FER's correspond to fading margins of about 10 and 5 dB, respectively. Also, note that $(1-q)^{-1}$ is the average length of frame error bursts (see Fig. 2). Since the Markov parameters p and q depend on the normalized Doppler bandwidth f_dT and the average FER P_E, we vary the burst error statistics in the simulations by choosing different values of P_E and f_dT.

The TCP and RLP throughput performance was obtained through simulations under different fading conditions. The simulation model considers only the data transfer phase, which dominates the steady-state performance. The procedures related to connection setup and teardown are not included in the simulation. The TCP layer program is taken from the TCP KA9Q source code [9]. The RLP implementation is as per IS-99 [7]. Every outgoing TCP packet is put into an interface buffer which is picked up by the RLP. The RLP breaks the TCP packet into 20-ms frames and sends them to the other end. As per IS-99, a storage buffer for resequencing out-of-sequence RLP data frames (128 frames worth buffer each on sending and receiving

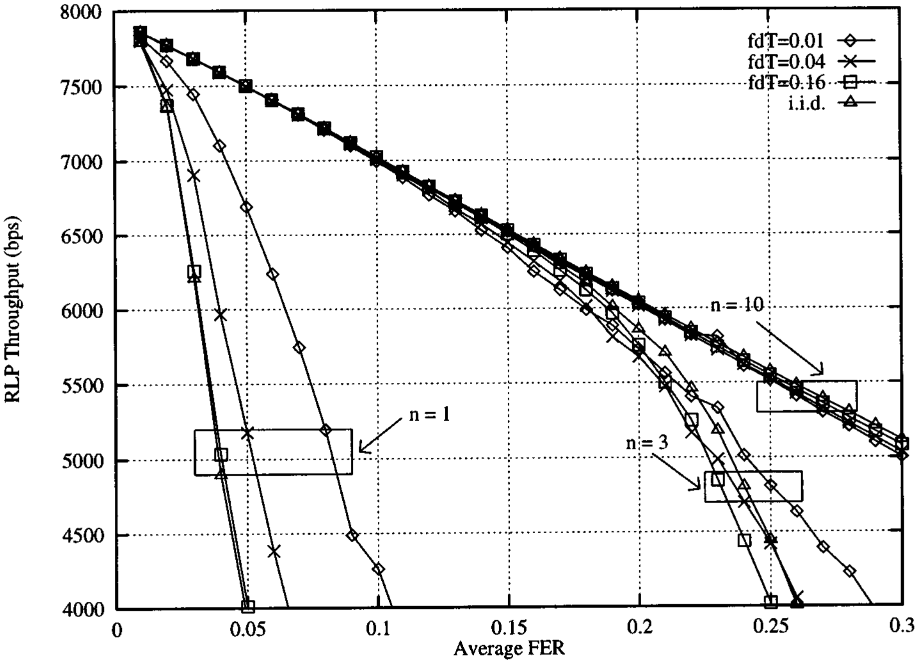

Fig. 3. RLP throughput versus average FER P_E at different values of normalized Doppler bandwidth f_dT and number of allowed retransmission attempts at the RLP layer n. TCP segment size $= 536$ bytes (i.e., default MSS).

direction) is used. In the event of a buffer overflow, for example, due to a long string of frame errors, a portion of buffer head (good and lost frames) is passed to higher layer to make room for the newly arriving frames. We consider two different TCP segment sizes namely, default MSS (536 bytes) and twice MSS (1072 bytes). It is further assumed that channel bandwidth is always available for the NAK frames sent by the receiving end, and that the NAK frames undergo no extra queueing delays caused by other traffic frames from the receiving end.

The RLP throughput (in bits/s) as a function of average FER (P_E) and normalized Doppler bandwidth (f_dT) is shown in Fig. 3 for a TCP segment size of 536 bytes. The various values of f_dT considered are 0.01, 0.04, and 0.16. At a carrier frequency of 900 MHz and $T = 20$ ms, the f_dT values of 0.01 and 0.16 correspond to mobile speeds of about 0.6 km/h (very slow moving user), and 9.6 km/h, respectively. From Fig. 2, when the fading margin is 5 dB (about 30% average FER), the f_dT values of 0.01, 0.04, and 0.16 represent average frame error burst lengths of about 26.4, 6.6, and 1.4 frames, respectively. The performance over an i.i.d. fading channel model is also plotted in Fig. 3 for comparison. The plots are parameterized by the number of retransmission attempts n that is allowed at the RLP layer to recover a lost frame. If RLP is unsuccessful sending a frame error free even after n retransmission attempts, it passes the recovery of that lost frame to the TCP layer. For the same set of conditions in Fig. 3, the throughput performance at the TCP layer is plotted in Fig. 4.

From Figs. 3 and 4, it can be seen that the throughput performance is poor if just one retransmission $(n = 1)$ is allowed at the RLP layer, compared to three or ten retransmissions ($n = 3$ or 10). The performance is consistently better with ten retransmissions compared to that with one or three retransmissions. This suggests that persisting at the RLP layer until the frame transmission succeeds is desired. However, over a range of average FER's from 1% to 15%, three retransmissions (which is the default value in the IS-99 standard) can achieve performances close to that of ten retransmissions. For voice applications, this FER range covers the typical operating region, beyond which a call, if not handed over, will be dropped due to excessive frame errors. In data applications, however, the link with a poor fading margin can be benefited by increased persistence at the RLP layer because the call need not be fully dropped as long as the realized throughput is adequate to support the application. For example, in applications like asynchronous data, flow control mechanisms (e.g., by use of V.42 between the base station and the remote PSTN DCE or software flow control using XON and XOFF) can throttle the source transmission if the channel throughput falls down below the source transmission rate. In this context, TCP throughputs achieved in excess of 4800 bps at $n = 10$, even when the average FER is as high as 30%, are significant. In other words, data applications can benefit by allowing more than three retransmissions at the RLP, particularly in the 20% to 30% FER range. As a side note, we point out that persistence at the RLP layer is beneficial on point-to-point circuit switched wireless links. However, it is not always beneficial in a "multiple access" reverse link scenario, particularly when the fading is very much correlated [13]. In fact, it is shown in [13] that in a multiple access scenario, longer persistence at the link layer would degrade the throughput performance under highly correlated channel fading conditions.

Another interesting observation in Figs. 3 and 4 is the effect of f_dT on the throughput for a given number of allowed RLP retransmissions. For large number of allowed retransmissions (e.g., $n = 10$), the throughput performance tends to become almost independent of f_dT (see the close bunching of curves for $n = 10$). Also, in such a case, the performance in i.i.d.

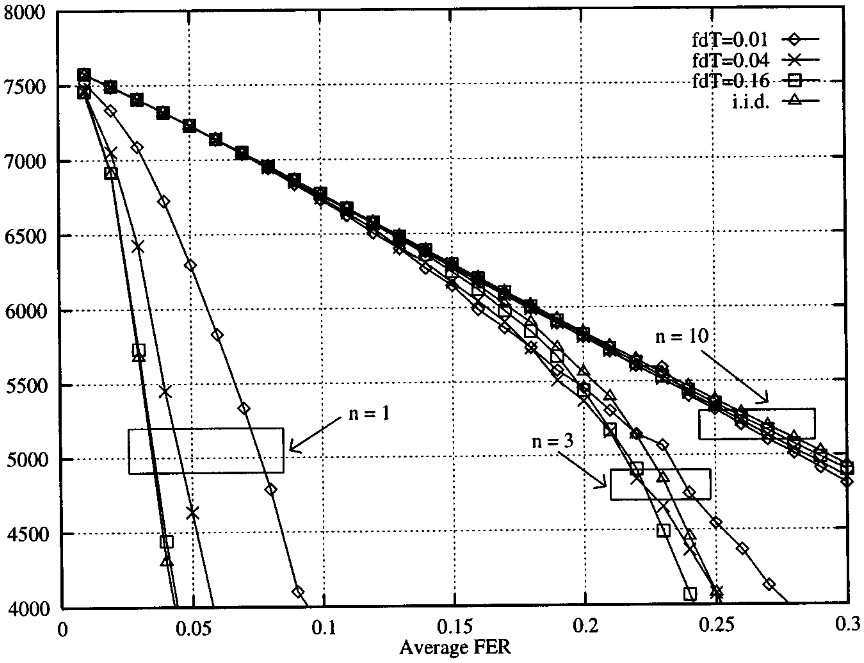

Fig. 4. TCP throughput versus average FER P_E at different values of normalized Doppler bandwidth f_dT and number of allowed retransmission attempts at the RLP layer n. TCP segment size = 536 bytes (i.e., default MSS).

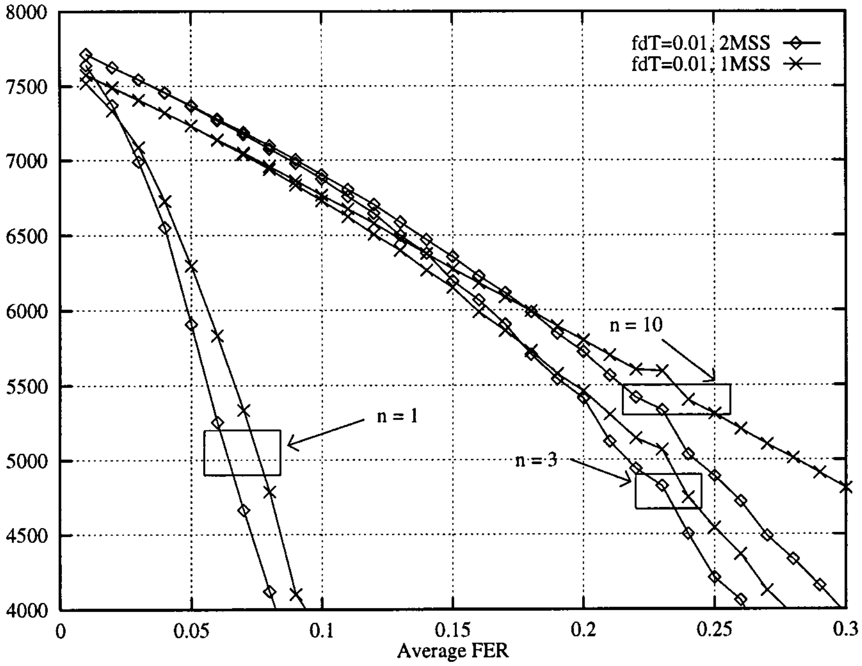

Fig. 5. TCP throughput versus average FER P_E at different TCP segment sizes (536 versus 1072 bytes). High correlation in fading (i.e., $f_dT = 0.0177$).

fading case remains best compared to correlated fading cases (although the performance difference is pretty small). However, as the number of allowed retransmissions is reduced (e.g., from $n = 10$ to $n = 3$), the performance difference between i.i.d. and correlated fading cases is clearly noticeable for average FER's higher than 15%. Performance crossovers between i.i.d. and correlated fading cases are found to occur at different average FER values. The crossover point varies depending on the value of f_dT (in other words, the degree of channel burstiness), and n (i.e., the amount of persistence at the RLP). For example, for $n = 3$ and $f_dT = 0.01$, the performance crossover occurs at 22% average FER such that i.i.d. case performance is better for $P_E < 0.22$, and $f_dT = 0.01$ case performance is better for $P_E > 0.22$. On the other hand, if n is reduced to one, the crossover point shifts to less than 1% average FER (not shown in the plots). Thus, in the case of $n = 1$, the i.i.d. case performance is worst compared to correlated fading cases for all FER's in the considered range of 1% to 30%. This is because not persisting long in correlated fading clearly avoids possible loss of consecutive retransmitted frames, which saves bandwith.

The effect of increasing TCP segment size from 536 bytes (default MSS) to 1072 bytes (2*MSS) on the TCP throughput performance is illustrated in Fig. 5 for $f_d T = 0.01$ (high correlation in fading). Here again, the plots are parameterized by different values of $n (=1, 3, 10)$. From Fig. 5, it is seen that in highly correlated fading, the TCP segment size of 2MSS (larger segment size) results in better performance than 1MSS (smaller segment size) at low FER's. This is because when link conditions are good, larger segment sizes improve throughput by sending segments successfully with less overhead compared to smaller segment sizes. However, when the link conditions are bad (i.e., high FER's), smaller segment sizes (1MSS) result in better performance than larger message sizes (2MSS). This is because when large segment sizes are used in bad link conditions, the TCP layer will witness increased segment losses and subsequent TCP retransmissions due to increased frame recovery failures at the RLP layer. The above performance crossover occurs at different average FER's depending on the value of n.

V. CONCLUSION

We presented the throughput performance of TCP/RLP protocol stack on correlated fading DS-CDMA wireless links. Specifically, we considered a scenario where a large data file is transferred from the base station host to a mobile terminal over a DS-CDMA wireless link. The dependence of the TCP/RLP throughput performance on the degree of channel burstiness and the amount of persistence at the RLP layer were evaluated and highlighted. It was shown that the burstiness in RLP frame errors, caused by the correlation in the multipath fading process, has a significant effect on the RLP and TCP throughput performance, particularly at low fading margins. On links with low fading margins, longer persistence at the RLP layer to recover lost RLP frames (more than the IS-99 specified three retransmission attempts at the RLP layer) was shown to result in increased TCP/RLP throughput. It was also shown that large TCP segment sizes perform better in good link conditions, whereas in bad link conditions small segment sizes perform better.

ACKNOWLEDGMENT

The authors would like to thank S. Veerepalli for his useful comments on this work.

REFERENCES

[1] W. R. Stevens, *TCP/IP Illustrated*. New York: Addison Wesley, 1994, vol. 1.

[2] T. V. Lakshman and U. Madhow, "The performance of TCP/IP for networks with high bandwidth-delay products and random losses," *IEEE/ACM Trans. Networking*, vol. 5, pp. 336–350, June 1997.

[3] H. Chaskar, T. V. Lakshman, and U. Madhow, "On the design of interfaces for TCP/IP over wireless," in *Proc. IEEE MILCOM'96*, San Diego, CA, 1996.

[4] H. Balakrishnan, V. N. Padmanabhan, S. Seshan, and R. H. Katz, "A comparison of mechanisms for improving TCP performance over wireless links," in *Proc. ACM Sigcomm'96*, CA, Aug. 1996.

[5] A. Kumar and J. Holtzman, "Comparative performance analysis of versions of TCP in a local network with a lossy link, part II: Rayleigh fading mobile radio link," Tech. Rep. WINLAB-TR-133, Nov. 1996.

[6] A. Chockalingam, M. Zorzi, and R. R. Rao, "Performance of TCP Reno on wireless fading links with memory," in *Proc. IEEE ICC'98*, vol. 2, GA, June 1998, pp. 595–600.

[7] *TIA/EIA/IS-99, Data services option standard for wideband spread spectrum digital cellular system*, 1995.

[8] *TIA/EIA/IS-95, Mobile station-base station compatibility standard for dual-mode wideband spread spectrum cellular system*, 1993.

[9] G. Bao, "Performance evaluation of TCP/RLP protocol stack over CDMA wireless link," *Wireless Networks*, pp. 229–237, 1996.

[10] W. C. Jakes, *Microwave Mobile Communications*. New York, NY: Wiley, 1974.

[11] H. S. Hang, "On verifying the first-order Markovian assumption for a Rayleigh fading channel model," in *Proc. IEEE ICUPC'94*, San Diego, CA, Sept. 1994, pp. 160–164.

[12] M. Zorzi, R. R. Rao, and L. B. Milstein, "On the accuracy of a first-order Markov model for data block transmission on fading channels," in *Proc. IEEE ICUPC'95*, Nov. 1995, pp. 211–215.

[13] A. Chockalingam, M. Zorzi, L. B. Milstein, and P. Venkataram, "Performance of a wireless access protocol on correlated Rayleigh fading channels with capture," *IEEE Trans. Commun*, vol. 46, pp. 644–655, May 1998.

A. Chockalingam (SM'98) was born in Rajapalayam, India. He received the B.E. (Honors) degree in electronics and communication engineering from the P.S.G. College of Technology, Coimbatore, India, in 1984, the M.Tech. degree with specialization in satellite communications from the Indian Institute of Technology, Kharagpur, India, in 1985, and the Ph.D. degree in electrical communication engineering from the Indian Institute of Science, Bangalore, India, in 1993. His Ph.D. dissertation work involved the design and performance evaluation of media access protocols for wireless networks.

During 1986–1993, he was with the Transmission R&D Division, Indian Telephone Industries, Ltd., Bangalore, India, where he was involved in a range of R&D projects including VSAT networks using TDMA and CDMA, forward-error correction, and variable-rate PSK modem implementations using ASIC and DSP technologies. From December 1993 to May 1996, he was a Post-Doctoral Fellow/Assistant Project Scientist with the Department of Electrical and Computer Engineering, University of California–San Diego, La Jolla, CA, where he conducted research in CDMA wireless communications. He was a Communication Systems Consultant to CommQuest Technologies, Inc., Encinitas, CA, during 1994–1995. Since May 1996, he has been with Qualcomm, Inc., San Diego, CA, where he is currently a Manager and Staff Engineer working in the area of system design and performance analysis of DS-CDMA wireless communication systems including Globalstar. His research interests lie in the area of wireless networks/protocols and digital communication systems.

Dr. Chockalingam was the recipient of the Communication Devices India, Ltd. (CDIL) Award for a paper published in the *IETE Journal*.

Gang Bao (M'93) was born in Hefei, China. He received the B.E. degree from the Huazhong University of Science and Technology, Wuhan, China, in 1984, the M.S. degree from the Institute of Automation, Chinese Academy of Sciences, Bejing, China, in 1987, and the Ph.D. degree in electrical and computer engineering from the University of Massachusetts, Amherst, in 1994.

During 1989–1994, he worked on stochastic optimizations and discrete event dynamic systems (DEDS) modeling, especially higher order derivative estimations in general Jackson-like queueing networks. He joined Qualcomm, Inc., San Diego, CA, in 1994, where he is currently a Staff Engineer working in the area of system design and performance analysis of DS-CDMA wireless communication systems. His major research interests are DEDS, wireless communications, and data networks.