Accurate approximation of ARQ packet delay statistics over Markov channels with finite round-trip delay

Michele Rossi, Leonardo Badia, Michele Zorzi
{mrossi,lbadia,mzorzi}@ing.unife.it
University of Ferrara, via Saragat 1, I-44100 Ferrara, Italy

***Abstract*— In this paper the packet delay statistics of a fully reliable Selective-Repeat ARQ scheme is investigated. The sender transmits packets whose error process is characterized by means of a two-state Discrete Time Markov Channel (DTMC). At the receiver these packets are checked for errors and ACK/NACK messages are sent back to the sender accordingly. It is assumed that the feedback message is known with no errors at the transmitter m channel slots (round-trip delay) after the packet transmission started. An appropriate Markov model has been previously developed in order to find the exact statistics of the delays experienced by ARQ packets. In this work, in order to reduce the computational complexity of such analysis, an approximate model is presented. The results obtained from the approximate approach are shown to be in excellent agreement with the ones derived from the exact analysis.**

I. Introduction

Reliable data transmission over error-prone channels requires mechanisms to recover from errors, which may occur with higher probability than the application can tolerate. In usual protocol stacks, error control is performed at multiple levels, e.g., at the physical layer by error correction codes, at the data-link layer by ARQ techniques, as well as at the transport layer by TCP. In order to cut the right tradeoff between data reliability, latency, and efficient bandwidth usage, error control techniques must be carefully designed and their performance well understood. The study of ARQ error control techniques in wireless systems has not enjoyed great popularity in recent years compared to, e.g., error correction coding strategies. This is mainly due to the type of application envisioned in these systems, i.e., voice and circuit-switched data, where strict delay guarantees are provided. With the extension of packet data and Internet services over wireless links, the increased delay tolerance of many applications and protocols leads to a paradigm shift, where error recovery by retransmission may be more efficient than protecting all data a priori by means of costly FEC.

The key point, when ARQ solutions are considered is that they directly interact with higher level, by determining both delay/jitter performance and error probability of higher level packets. For these reasons, their correct configuration is key in achieving the needed higher level QoS. Hence, an accurate study of the delivery delay process at the ARQ level is pivotal in order to understand the interaction between the higher level performance and the link layer retransmission process.

In ARQ, the transmitter sends packets (PDUs) consisting of payload and error detection codes. At the receiver side, based on the outcome of the error detection procedure, acknowledgment messages are sent back to the transmitter (ACK or NACK, according to the result of error detection). The sender performs packet retransmissions based on such acknowledgments. In general, ARQ protocols are variants of the following basic schemes: stop-and-wait (SW), go-back-N (GBN) and selective repeat (SR). In this paper we consider the SR scheme (the most efficient), where packets are transmitted continuously, and only negatively acknowledged packets are retransmitted, i.e., retransmissions are selectively triggered by NACK messages. We observe that, when the round-trip delay goes to zero all the mentioned schemes become identical. According to [1] and [2] we refer to this situation as *ideal SR* ARQ scheme.

In the presence of the ARQ protocol, we can subdivide the overall PDU delay in three contributions. The first is due to the queueing delay in the source buffer, i.e., the time between the PDU release by higher levels and the instant of its first transmission over the channel. The second contribution is the time between the first transmission and the correct reception of the PDU (only depends on the channel behavior). The last delay is due to the time spent in the receiver re-sequencing buffer. In fact, even if the sender transmits packets in order, due to random errors and consequent retransmissions, they can arrive out of sequence. Hence, correctly received PDU with higher identifier must wait in the receiver re-sequencing buffer until all the PDUs with lower identifier have been correctly received. This last term is the most complicated because, by considering a tagged PDU, it depends on errors experienced by all PDUs sent in the same round-trip in which the tagged one has been transmitted for the first time. These quantities will be referred to as *queueing delay*, *transmission delay* and *re-sequencing delay*, as usually done in the literature [2]. In addition, we define *delivery delay* as the sum of the second and third term. In this work we focus on the computation of this last term.

Several studies have been performed on the delay performance of the SR protocol over a wireless channel [1][2][6][7][8][9][10]. In [6] Konheim derived the exact distribution of PDU delays with finite round-trip delay, but considering an *i.i.d.* error process. Rosberg and Sidi in [9] analyzed the joint distribution of transmitter and receiver buffer occupancies over a static channel. In [10] Zorzi and Rao, considered the *ideal SR* scheme, and proved the effectiveness of Markov model by means of comparison with a simulated fading process. The time varying channel has been investigated for the first time by Fantacci in [1], where the *ideal SR* scheme is considered. This approach leads to an analytical lower bound on delay performance with respect to the situation of a finite round-trip delay. In [2] Kim and Krunz accounted for a time varying channel, a finite round-trip delay and a Markovian traffic source. Here, a mean analysis is developed for all the ARQ delay contributions, in the computation of the source queueing delay the *ideal SR* is considered and the mean re-sequencing delay is obtained from an approximate approach and by considering *Heavy traffic* condition.

This work has been partially supported by ERICSSON Research.

We study the delay performance of a fully reliable SR ARQ scheme, in the presence of a time varying channel [3][4][5] and of a finite round-trip time. In [13] an exact analysis for the computation of the delivery delay statistics of SR-ARQ packets is presented, which accounts for a finite round trip delay and for a time varying channel. The complexity of such approach grows exponentially with the round trip time value. In this paper, an approximate approach with very limited complexity (linear in the round-trip delay) is developed in order to find the delivery delay statistics of SR-ARQ packets; its accuracy is demonstrated by comparison with the results in [13].

The remaining part of the paper is organized as follows: in Section II the ARQ policy and the channel model are described, in Section III the approximate approach for the computation of the delivery delay statistics is presented, in Section IV results are reported and finally, in Section V, some conclusions are given.

II. Model for ARQ queueing and transmission processes

We consider a pair of nodes, say a transmitter and a receiver, that communicate data packets through a noisy wireless link and use a fully reliable Link Layer protocol (unlimited retransmission attempts) to counteract channel impairments. In the forward direction data packets (ARQ PDUs) flow, while in the backward one only ACKs and NACKs flow (ACKs and NACKs are assumed error free). Moreover, we assume that both transmitter and receiver have unlimited buffer size and they adopt the following Selective-Repeat ARQ protocol (a generalization of the protocol described in [11]) at the Link Layer.

The sender continuously transmits new PDUs in increasing numerical order as long as ACKs are received. After each PDU reception, the receiver checks for packet errors and replies with an ACK/NACK accordingly. When the generic PDU is transmitted, say PDU i, the sender must await an ACK message for that packet until after it finishes the transmission of up to $m-1$ subsequent PDUs (new or retransmitted), $1 \leq m \leq \infty$, where m is the round-trip delay (m is commonly referred to in the literature [12] as the *ARQ window size*). At this point, PDU i must be retransmitted (in the case of NACK) or a new PDU has to be sent (in the case of ACK). The time is slotted and the slot time corresponds to a single PDU transmission. The wireless channel is characterized by means of a two-state Discrete Time Markov Chain (DTMC). Let 0 and 1 denote successful and erroneous PDU transmission in a given slot, respectively, and let:

$$\mathbf{P} = \begin{pmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{pmatrix} \tag{1}$$

be the channel transition probability matrix. The steady-state channel error probability is given by:

$$\varepsilon = \frac{p_{01}}{p_{10} + p_{01}} \tag{2}$$

while the average error burst length is given by $b = \frac{1}{p_{10}}$. We define as $\mathbf{P}(i)$ the i-step transition probability matrix, computed as follows:

$$\mathbf{P}(i) = \mathbf{P}^i = \begin{pmatrix} p_{00}(i) & p_{01}(i) \\ p_{10}(i) & p_{11}(i) \end{pmatrix} \tag{3}$$

Moreover, we consider a *Heavy Traffic* condition, i.e., once a PDU is correctly transmitted, a new one is always present in the source buffer. This assumption is justified by taking into account a TCP file transfer (FTP-like session or video/audio continuous data streaming) as packet source at the transmitter. The reliable ARQ completely avoids TCP timeouts (when the channel error is not too large) and the TCP level, after filling the bandwidth-delay product, behaves as a continuous packet source (the TCP window size is not decreasing because error recovery is never triggered).

As a last observation, where the *Heavy Traffic* assumption is not verified, the delivery delay computed with our model is an analytical upper bound. Also in this case, our approach is useful as a worst case analysis.

III. Approximation of the delivery delay statistics

In order to compute the statistics for a single PDU transmitted usding Selective Repeat ARQ, we develop a model which tracks the successful delivery state of the PDU of interest (referred as *tagged PDU*), as well as all previous PDUs. Suppose that the tagged PDU is transmitted for the first time in slot $t = m$. This implies that all previous PDUs (i.e., those whose identifier is smaller than the tagged PDU) excluding the $m-1$ PDUs transmitted in slots 1 through $m-1$ have been successfully received, and that in slot 0 a successfull transmission occurred; in any other case in slot m we would have a retransmission. In more detail, in Selective Repeat ARQ when a packet is transmitted erroneously at time t it is always scheduled for retransmission m slots apart (time $t+m$), only after it successfull transmission the slot it occupies is released and can be used for a new packet transmission. Therefore, a new transmission in a given slot implies that a correct transmission occurred m slots earlier. The tagged PDU is finally released upon correct reception of all PDUs transmitted in slots 1 through m. Note also that all PDUs transmitted for the first time during slot $t > m$ must have a larger id that the tagged PDU, and therefore do not affect its delivery process. We can then ignore all future PDU arrivals in our study.

The problem to be solved is therefore to find the time it takes for all PDUs transmitted in slots 1 through m to be eventually received correctly, given that a correct PDU transmission occurred in slot 0. The statistics associated with this time quantity is called delivery delay statistics. Consider the evolution of the system after slot m. If slot 1 contained an erroneous transmission, a retransmission will be scheduled in slot $m+1$. If this retransmission is successfull, then slot $m+1$ will be marked as *resolved*, otherwise it will be marked as *unresolved* (with the effect of a further retransmission in slot $2m+1$ and so on until success). In [13] an exact analysis has been presented, whose main drawback is that its complexity grows exponentially with m (in this analysis the state of each slot 1 through m is tracked until all these slots are resolved). Hence, for large values of the round trip delay, the computation of the exact statistics becomes both memory and time expensive.

In this Section, we propose a simple approximate approach that allows to reduce the computational complexity enabling the computation of the statistics for large m. In more detail, instead of tracking the exact position of each unresolved slot, we propose to analyze the resolving process simply by tracking the number of remaining slots to be resolved.

To do that, we subdivide the time in rounds of m slots and we build a Markov chain in which the state is represented by $X'(r) = (n_e(r), C(r))$, $r \geq 1$ where $n_e(r)$ and $C(r)$ are the number of unresolved slots in round r and the channel state in the last slot of round r, respectively. This Markov chain evolves round by round, i.e., m slot at a time. Round 1 corresponds to the first round, i.e., the one comprising slots 1 through m. As mentioned above, we consider that the *tagged PDU* is transmitted for the first time at time m, and that at time 0 the channel is constrained to be in the good state, otherwise the slot in position m would be occupied by a retransmission.

Note that, in this Markov chain, the knowledge about the position of unresolved slots is completely ignored. Moreover, in order to keep track of errors and corrections, we arbitrarily assume that all unresolved slots are deterministically placed at the end of each round. This simple approximation rule allows us to neglect the exact position of each unresolved slot and to achieve a computational complexity linear in m.

Let $\phi_{ij}(k,n), i,j \in \{0,1\}$ be the probability that there are k successful slots in $\{0,1,\ldots,n-1\}$ and that the channel state is j at time n, given that the channel state was i at time 0. $\phi_{ij}(k,n)$ can be computed recursively as follows (see [14]):

$$\begin{aligned}\phi_{ij}(k,n) &= \phi_{i0}(k-1,n-1)p_{0j} + \phi_{i1}(k,n-1)p_{1j} + \\ &+ \delta_{i,j}\delta(k)\delta(n)\end{aligned} \tag{4}$$

where $\phi_{ij}(k,n) = 0$ for negative values of either k or n, $\delta_{ij} = 1$ if $i = j$ and zero otherwise, and $\delta(k) = \delta_{k,0}$. A closed form for that function can be found in [15]. Moreover, let $\varphi_{ij}(k,n)$ be the probability that there are k successful slots in $\{1,\ldots,n\}$ and that the channel state is j at time n, given that the channel state was i at time 0. It is straightforward to relate $\varphi_{ij}(\cdot,\cdot)$ to $\phi_{ij}(\cdot,\cdot)$ by noting that $\varphi_{10}(j,n) = \phi_{10}(j-1,n)$ and $\varphi_{01}(j,n) = \phi_{01}(j+1,n)$, whereas in the other two cases ($i=j=0$ and $i=j=1$) they are the same.

Now, let $\Psi_{0j}(e,r)$ be the probability to have e ($1 \le e \le m$) unresolved slots in round r and that the channel in the last slot of round r is j given that the channel in slot 0 was correct. This function can be computed in a recursive way as follows:

$$\Psi_{0j}(e,r) = \begin{cases} \varphi_{0j}(m-e,m) & r=1 \\ \displaystyle\sum_{k=e}^{m}\sum_{i\in\{0,1\}} \Psi_{0i}(k,r-1)\mathcal{R}_{ij}(k-e,k) & r>1 \end{cases} \tag{5}$$

where,

$$\mathcal{R}_{ij}(q,k) = \begin{cases} \varphi_{ij}(q,k) & k=m \\ \displaystyle\sum_{c\in\{0,1\}} p_{ic}(m-k)\varphi_{cj}(q,k) & k<m \end{cases} \tag{6}$$

where the function $\mathcal{R}_{ij}(q,k)$ is used to compute the probability of resolving q slots over k in round r, and that the channel state in the last slot or round r is j given that the channel state in the last slot of the previous round ($r-1$) was i, and that all unresolved slots are deterministically grouped at the end of each round. Note that the recursive expression (Eq. (5)) is initialized (round $r=1$) by exploiting the knowledge of the channel at time 0 and computing the mean probability to have e erroneous transmissions[1] in any order in that round. Moreover, the probability to have e unresolved slots, $1 \le e \le m$, at the end of round r, $r > 1$ ($\Psi_{0j}(e,r)$) is obtained by considering the probability to have k ($\Psi_{0i}(k,r-1), e \le k \le m$) unresolved slots after $r-1$ rounds and that exactly $k-e$ of these slots are resolved in round r ($\mathcal{R}_{ij}(k-e,k)$). Finally, these probabilities are summed over $e \le k \le m$ and over $i \in \{0,1\}$ to account for all the admitted values of unresolved slots and channel states at the end of round $r-1$. Observe that the probability to be in state $X'(r) = X' = (n_e, C)$, $1 \le n_e \le m$, $C \in \{0,1\}$, $r \ge 1$ is given by $\Psi_{0C}(n_e,r)$.

Now, let us write the time index t as $t = \xi m + \eta$, where $\xi \ge 0$ and $1 \le \eta \le m$ are the number of full rounds and the number of slots in the current round covered by t, respectively. With this decomposition, the current round is round $\xi+1$, whereas the previous one is round ξ.

In the following, the approximate probability to have a delay equal to t is computed in two ways, depending on the assumption made about the position of the last unresolved slot in round ξ. More in detail, the number of slots, e ($1 \le e \le m$), that are yet to be resolved at the end of round ξ is computed by means of the function $\Psi_{0j}(e,\xi)$. These unresolved slots are considered to be grouped in a single burst, i.e., there are no isolated unresolved positions. In the first approach, the last unresolved slot of such burst is always considered to be in position $(\xi+1)m$, i.e., all unresolved slots are deterministically constrained to be at the end of the round. In the sequel, we refer to this correcting strategy as *Burst at the End (BE)*. In the second case, instead, we admit that the burst of unresolved slots can be cyclically shifted. This strategy is called *Shifted Burst (SB)*.

Before starting with the detailed description of the approaches just introduced, let us define the function ρ as the probability that the channel in the last slot of round ξ is j and that exactly e ($1 \le e \le m$) slots are yet to be resolved at the end of round ξ and that these e slots are all resolved in round $\xi+1$ (current round) given that the last unresolved slot is in position x, $e \le x \le m$. Formally:

$$\rho(j,e,\xi|x) = \Psi_{0j}(e,\xi)p_{j0}(x-e+1)p_{00}^{e-1} \tag{7}$$

This function is used in what follows to compute the probability of resolving a burst accounting for its length (e), the position occupied by its last element (x) and the channel state at the end of the current round (j). In the following we report the detailed description of the two approaches:

- (BE) The delivery delay statistics is computed in the following way:

  $$P_d[t] = \begin{cases} \displaystyle\frac{1}{m}\sum_{j\in\{0,1\}}\sum_{e=1}^{m}\rho(j,e,\xi|m)u[\eta-e] & \eta \ne m \\ \displaystyle\frac{1}{m}\sum_{j\in\{0,1\}}\sum_{e=1}^{m}\rho(j,e,\xi|m)e & \eta = m \end{cases} \tag{8}$$

  where $t = \xi m + \eta$, with $\xi \ge 0$ and $1 \le \eta \le m$. The function $u[\cdot]$ is defined as follows, $u[n] = 1$ for $n \ge 0$, whereas $u[n] = 0$ for $n < 0$, n integer.
  Specifically, the probability to have the first passage through states $(0,C)$ in round $\xi+1$, i.e., $t \in \{\xi m + 1, \ldots, \xi m + m\}$ is computed by considering all the unresolved slots inherited from round ξ to be placed at the end of round $\xi+1$. Moreover, this probability is subdivided among slots $\xi m + \eta, \forall\, \eta \in \{1,\ldots,m\}$ by assuming that the final unresolved slot is distributed by means of the function $\mathcal{P}(x|e)$ that is the approximate[2] probability that the burst ends in position x given that it consists of e slots. $\mathcal{P}(x|e)$ is given as follows:

[1] Note that each erroneous transmission through slot 1 to m corresponds to an unresolved slot.

[2] The approximation consists in the assumption that each slot has the same probability to be the final one.

$$\mathcal{P}(x|e) = \begin{cases} 0 & x < e \\ 1/m & e \le x < m \\ e/m & x = m \end{cases} \quad (9)$$

To sum up, in this approach we first resolve all the unresolved slots inherited from the previous round by deterministically grouping it at the end of the round (the last unresolved slot is always in position m of round $\xi + 1$). By this way we obtain the probability to resolve the initial window in any position of the current round. After that, we subdivide such probability among slots in the current round by considering the last unresolved slot to be uniformly distributed between position 1 through m.

- (SB) In this approach, we also consider that the unresolved slots inherited from round ξ are grouped in a single burst, but unlike in BE, we resolve such burst by considering that the position of its last element is uniformly distributed (with probability $1/m$) between position 1 through m of the current round. With this assumption the delivery delay statistics can be written as:

$$P_d[t] = \begin{cases} \dfrac{1}{m} \displaystyle\sum_{j\in\{0,1\}} \sum_{e=1}^{m} \rho(j,e,\xi|\eta) u[\eta - e] & \eta \neq m \\ \dfrac{1}{m} \displaystyle\sum_{j\in\{0,1\}} \sum_{e=1}^{m} \Big[\rho(j,e,\xi|m) + \\ \displaystyle\sum_{y<e} \rho'(j,e,\xi|y)\Big] & \eta = m \end{cases} \quad (10)$$

where,

$$\rho'(j,e,\xi|y) = \Psi_{0j}(e,\xi) p_{j0} p_{00}^{y-1} p_{00}(m - e + 1) p_{00}^{e-y-1} \quad (11)$$

With Eq. (10), we compute, for each $\eta \in \{1,\ldots,m\}$ in round $\xi + 1$, the probability to resolve the burst given that its last slot is in position η, where $\eta \ge e$. Thus, when $\eta \ge e$, we guarantee that the first slot of the burst is in position f with $f \ge 1$. In this case, we correct the burst of unresolved slots by means of the function $\rho(j,e,\xi|x)$ by letting[3] $x = \eta$.

In the case $\eta = m$ instead, the contributions of m-cyclically shifted versions of the burst are also considered. In more detail, we split the e-sized burst in two parts where the first part is composed by the y slots in position $\{1,\ldots,y\}$, whereas the second part is composed by the $e-y$ slots in position $\{m-e+y+1,\ldots,m\}$. These contributions are resolved by means of the function $\rho'(j,e,\xi|y)$. ρ' is used to evaluate the probability that the channel in the last slot of round ξ is j and that exactly e ($1 \le e \le m$) slots are yet to be resolved at the end of round ξ and that these e slots are all resolved in round $\xi+1$ (current round) given that they are subdivided in two bursts, where the first one occupies positions $\{1,\ldots,y\}$ and the second occupies positions $\{m-e+y+1,\ldots,m\}$.

To sum up, in both presented approaches, we consider that the slots that are yet to be resolved are grouped in a single burst. In the BE approach, we first compute the resolving probability considering the burst deterministically placed at the end of the round and then we subdivide this contribution among the possible positions in a uniform way. In the SB approach, instead, the last slot of the burst is considered as uniformly placed (with probability $1/m$) directly at the evaluation of the resolving probability.

[3]The same function has been used in the BE approach, but with $x = m$.

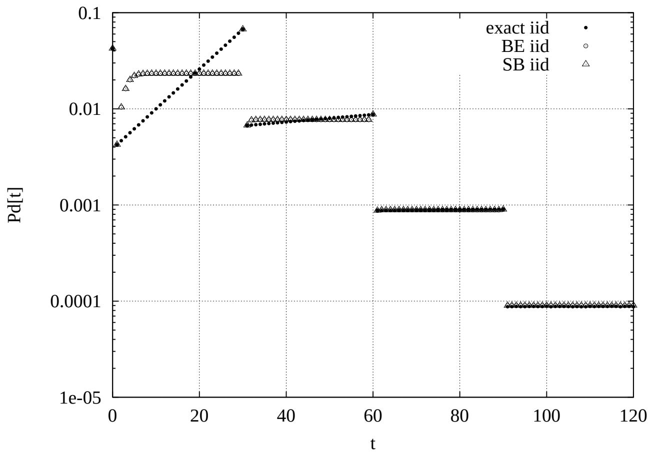

Fig. 1. Comparison between approximation and exact statistics, $m = 30$, $\varepsilon = 0.1$, *iid* channel.

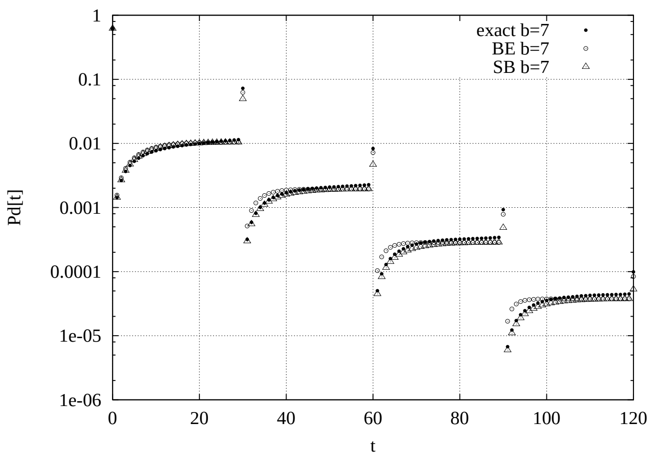

Fig. 2. Comparison between approximation and exact statistics, $m = 30$, $\varepsilon = 0.1$, correlated channel with $b = 7$.

IV. Results

In order to better clarify the results presented in this Section, let us remember that what we call $P_d[t]$ is the probability to resolve the initial window (PDUs transmitted in slot 1 through m) in a given number of slots ($k \ge 0$). The tagged packet delivery delay, instead, is the number of slots elapsed between the first transmission of the tagged packet and the instant in which it is released by the receiver re-sequencing buffer. The two delays above only differ for the sum of path delay and physical layer processing (that are both constant terms). In the following, we will refer to $P_d[t]$ as the tagged packet delivery delay statistics by keeping in mind this difference.

The delivery delay statistics ($P_d[t]$) has been computed according to the analysis presented in [13], for various values of the channel error probability, ε, and channel burstiness, b. Let us focus on the comparison between the exact analysis and the

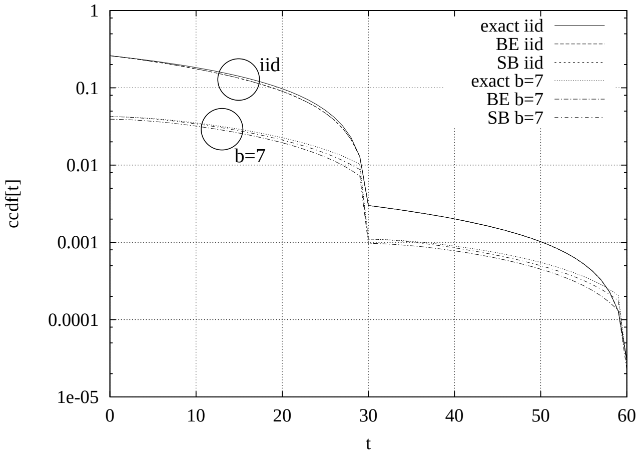

Fig. 3. Complementary cumulative delivery delay distribution, $ccdf[t]$: comparison between approximate and exact case, $m = 30$, $\varepsilon = 0.01$.

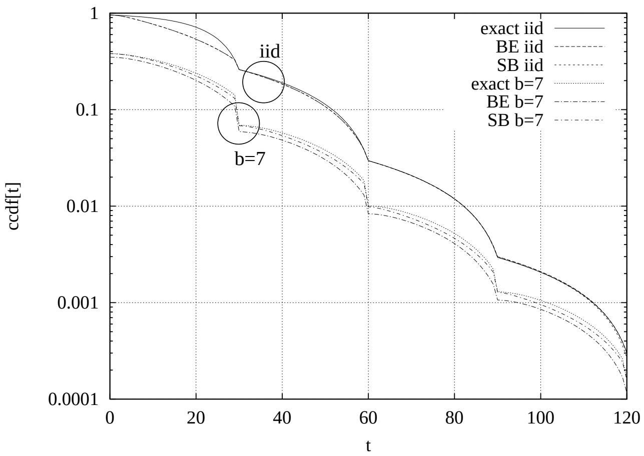

Fig. 4. Complementary cumulative delivery delay distribution, $ccdf[t]$: comparison between approximate and exact case, $m = 30$, $\varepsilon = 0.1$.

approximate approaches presented in Section III. This comparison is reported in Figs. 1 and 2 for an *iid* and a correlated channel ($\bar{b} = 7$), respectively. In general, the approximate approaches are in good agreement with the exact statistics.

The only region where the approximations fail is for uncorrelated channel and at low delays. In an *iid* channel, in fact, errors do not occur in burst, and so the approximation made in Section III that unresolved slots are disposed in a bursty way in this case does not hold. However, the effect of this approximation vanishes very quickly as the delay increases and the approximation becomes very close to the exact curve. In the *iid* case, both BE and SB approaches give the same results. When the channel is correlated (Fig. 2), instead, the statistics obtained from BE and SB are in good agreement with the exact curve for any value of the delay. The BE approach overestimates the delivery delay statistics at the beginning of each round. The SB approach, instead, appears to underestimate the exact statistics for any value of the delay. Moreover, the estimate obtained from BE degrades as the delay (t) increases until, for a very large t, all points in the round are aligned over a straight line. SB, instead, gives a good approximation also for large values of t. The points derived using SB are the closest to the exact curve, except for $t = im$, $i \geq 1$, where the best estimate is given by BE. For what concerns the complementary cumulative delay distribution ($ccdf[t]$), (Figs. 3-4) BE gives the best estimate for any value of b. Moreover, in the independent case, the burst assumption only results in a small discrepancy regarding the first round. For any other value of t, exact statistics and approximation match almost perfectly. In the correlated case ($\bar{b} = 7$), the BE approach gives the best complementary cumulative delay distribution estimate. From the obtained results, we can conclude that the estimate of the delivery delay statistics is reasonably accurate for any b, so approximate methods could effectively be used in real systems enabling a fast and less memory expensive computation of delay distributions also for large values of m.

V. Conclusions

In this paper we studied the delivery delay performance of a Selective Repeat ARQ scheme over a two-state Discrete Time Markov Chain. By presenting a simple approximate approach, we successfully avoid a burdensome analysis (whose complexity depends exponentially on the round trip delay). The statistics obtained using this approximate analysis are in excellent agreement with exact curves while keeping the complexity linear in m. The approximate distributions and their main characteristics are compared for several values of the channel error probability and error correlation against exact curves.

References

[1] R. Fantacci, "Queueing analysis of the selective repeat automatic repeat request protocol for wireless packet networks," IEEE Trans. Veh. Technol., vol. 45, pp. 258-264, May 1996.

[2] J. G. Kim, M. M. Krunz, "Delay analysis of Selective Repeat ARQ for a Markovian Source Over a Wireless Channel", IEEE Trans. Veh. Tech., vol. 49, pp. 1968-1981, Sept. 2000

[3] M. Zorzi, R.R. Rao, L.B. Milstein, "Error statistics in data transmission over fading channels," IEEE Trans. Comm., vol. 46, pp. 1468-77, November 1998.

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

[5] M. Zorzi, R.R. Rao, "On the Statistics of Block Errors in Bursty Channels," IEEE Trans. Comm., vol. 45, pp. 660-667, June 1997.

[6] A. G. Konheim, "A queueing analysis of two ARQ protocols," IEEE Trans. Commun. vol. 28, pp. 1004-1014, July 1980.

[7] M.E. Anagnostou and E.N. Protonotarios, "Performance analysis of the Selective Repeat ARQ protocol," IEEE Trans. Comm. vol. 34, no. 2, pp. 127-135, February 1986.

[8] Z. Rosberg and N. Shacham, "Resequencing delay and buffer occupancy under the Selective Repeat ARQ," IEEE Trans. Inform. Theory, vol. 35, pp. 166-173, January 1989.

[9] Z. Rosberg, M. Sidi, "Selective-Repeat ARQ: The joint distribution of the transmitter and the receiver resequencing buffer occupancies," IEEE Trans. Commun., vol. 38, pp. 1430-1438, September 1990.

[10] M. Zorzi, R.R. Rao, "On channel modeling for delay analysis of packet communications over wireless links," 36th Annual Allerton Conference, Allerton House, Monticello, IL, September 1998.

[11] S. Lin, D.J. Costello, M.J. Miller, "Automatic-repeat-request error control schemes," IEEE Comm. Mag., vol. 22, n. 12, pp. 5-17, December 1984.

[12] D. Bertsekas, R. Gallager, "Data Network," Prentice Hall, 1992.

[13] M.Rossi, L.Badia, M.Zorzi, "Exact delay statistics of ARQ packet delivery delay over Markov channels with finite round-trip delay," submitted for publication to IEEE GLOBECOM 2003.

[14] R. A. Howard, "Dynamic Probabilistic Systems," New York: Wiley, 1971.

[15] M. Zorzi, R.R. Rao, "Lateness probability of a retransmission scheme for error control on a two-state Markov channel," IEEE Trans. Commun., vol. 47, pp. 1537-1548, October 1999.

[16] Sheldon M. Ross, "Stochastic Processes, 2nd Edition," John Wiley & Sons 1996.