

RECEIVER BUFFER BEHAVIOR FOR THE SELECTIVE REPEAT PROTOCOL OVER A WIRELESS CHANNEL: AN EXACT AND LARGE-DEVIATIONS ANALYSIS

KOEN DE TURCK AND SABINE WITTEVRONGEL

SMACS Research Group, Department TELIN (IR07), Ghent University
St.-Pietersnieuwstraat 41, 9000 Gent, Belgium

ABSTRACT. In this paper, we formulate and analyze a model of the resequencing buffer at the receiver's side for the Selective Repeat protocol over a general class of transmission channels. Thanks to its efficiency, Selective Repeat is a ubiquitous error control mechanism in many different settings, in particular in wireless protocols such as WiMax and WiFi.

In view of the correlated nature of transmission errors over wireless channels, the receiver buffer model considers a general Markovian error process. We provide both an exact mathematical analysis of the receiver buffer behavior as well as a computationally efficient large-deviations result. An asymptotic analysis of the delay is also given. Numerical examples show that the correlation of the error process has an important influence on the performance of the receiver buffer.

1. Introduction. Retransmission protocols are a crucial part of almost any means of telecommunications, in almost any layer of the protocol stack. They play an especially important role in protocols over wireless links such as WiFi and WiMAX, as these links often exhibit an error process that is both bursty and has high error rates. The principle is simple: the sender transfers data (in the form of a numbered packet stream) to a receiving unit over an unreliable channel. The receiver acknowledges the status of each of the received packets, to which the sender responds by resending the incorrectly received packets. Many retransmission protocols have been proposed, each distinguished by the way that the retransmissions and the acknowledgements are organized.

Of the basic types, the Selective Repeat protocol is the most efficient protocol in terms of throughput, and as such it is widely implemented in numerous appliances. However this efficiency comes at a price, as it is one of the most complex retransmission protocols to implement and to analyze.

As packets might arrive out of order, Selective Repeat requires – in contrast to other protocols – a resequencing buffer at the receiver's side, in which a correctly received packet can be stored until the original packet order can be restored.

2000 *Mathematics Subject Classification.* Primary: 68M20, 60K25; Secondary: 90B22.

Key words and phrases. performance analysis, wireless communications, queueing theory, large deviations, retransmission protocols.

This paper was presented at the QTNA2009 conference which was held in Singapore during July 29–August 1, 2009. An earlier and brief version of this paper was published in the Conference Proceedings. The reviewing process of the paper was handled by Wuyi Yue, Yutaka Takahashi and Hideaki Takagi as Guest Editors.

We provide in this paper both an exact and a large-deviations analysis of the receiver buffer behavior. The latter results in computationally simple and insightful formulas of the essential performance characteristics. It is our belief that these may be more useful to the network analyst than an exact and elaborate analysis that is tedious to compute and to interpret. The analysis in this paper might prove to be instructive beyond the present application, as it shows a way to obtain efficiently computable performance measures for a model that is particularly struck by the curse of dimensionality. That is, the state space explodes exponentially as a function of the roundtrip time, which means that traditional solution techniques become intractable for even moderate values of that parameter. We show in this paper that we can nevertheless extract with only minor efforts the decay rate of the buffer content distribution, a performance measure that is often argued to be among the most useful for teletraffic engineers.

The goal of the paper is first of all to study the performance characteristics of the SR protocol and especially to clarify the behavior of the reordering buffer at the receiver side and the true effect of error correlation. These issues have not yet been properly clarified before. A further goal of the paper is to communicate another successful application of large-deviation theory to the performance modeling community.

Previous work has been done on the queueing performance of the Selective Repeat protocol, most papers focusing on the buffer at the sender's side. The classic paper by Konheim [17] provides an exact solution, that however suffers from the above-mentioned curse of dimensionality. Various approximations with diverse levels of accuracy and mathematical justification have been proposed in [16, 11, 1]. More recent papers also study the impact of correlation on the protocol performance (see e.g. [8, 7, 4, 13]). A key asset of the present paper is that we provide mathematical proofs of the proposed approximations.

Papers which focus on the queue at the receiver's side have been less numerous. In the papers [5, 19], independent and identically distributed (IID) error processes were considered, whereas in the present paper we extend this to general Markovian error processes. More recent efforts are reported in [6] and [22]. In [6], as in the present work, Markovian instead of IID processes were considered, however in the arrival process instead of the error process. The analysis of [22] shares with this paper a strong emphasis on tractable formulas, but it is applied to a situation where packet delays are variable and packet errors are non-existent, quite different from the lower-layer Selective Repeat setting that is the subject of this paper.

The outline of this paper is as follows. In section 2, we describe the mathematical model of the resequencing buffer. Section 3 contains a method to compute the exact solution. The large-deviations analysis is performed in section 4, and in section 5 we look at how exactly large buffer occupancies occur. We study the asymptotic distribution of the packet delay in section 6. Thereafter, in section 7, we explore the similarity with a well-known combinatorial problem, the coupon-collector problem. Next, we show some numerical results in section 8. We conclude in section 9.

2. Mathematical model of the resequencing buffer. We extend the mathematical model of [5] and [19] to more general error processes. The sender is assumed to operate under saturated-buffer conditions, which means that there are always packets available to be sent. This assumption frees us from specifying any arrival process, and can moreover be considered as a worst case. Indeed, the fact

that sometimes there are no new packets to be sent, can only have a positive effect on the performance of the system. We assume that the time is divided into slices of equal length, called slots, and the transmission of one packet is assumed to take exactly one slot. The receiver sends an acknowledgement message of each packet. If a packet is incorrectly received, the transmitter sends a new copy, that arrives at the receiver exactly L slots later. This duration of L slots is called the roundtrip time or the feedback delay of the channel, and has a tremendous effect on the performance of the system, as we will see. If the transmitter receives a positive acknowledgement message, it just takes a fresh packet and attempts to send it.

When looking at a sample path as in Figure 1, it is handy to fit it into a rectangular strip of width L . Time proceeds as in a Western text, left to right, top to bottom. As in [5] we call a vertical line in such a plot a 'group' and note that every transmission attempt of the same packet takes place during the same group. How can we determine how many packets are waiting in the resequencing buffer at an arbitrary slot boundary t ? It is simply the number of all the correctly received packets that were first transmitted after the eldest erroneously received packet (EEP). This EEP can be detected as the longest uninterrupted gray region looking from bottom to top. The buffer content at time t is hence simply the number of white squares, counted from time t backwards until we meet the EEP (the dashed squares do not count as they represent packets that are elder than the EEP). This is basically the reasoning in [5] and in [19].




FIGURE 1. First method to determine the resequencing buffer content at a given time instant ($L = 6$). Erroneous transmissions are indicated in gray. Dashed boxes represent transmissions of packets that have already left the buffer. We identify the EEP (here: packet 10) and then count the number of packets with higher sequence numbers that are already successfully transmitted. We thus find 10 packets in the resequencing buffer.

In the present paper, we take a related, and arguably a somewhat more elegant approach (and in case of correlated errors, perhaps the only feasible one). We define a process that runs backwards in time from a random slot boundary onwards, and we count the number of correct packets that we encounter, as illustrated in Figure 2. We stop the process when there has been a correct packet in every group. It appears that the buffer content at that random slot boundary is the total number of correct

packets minus L . This means that the buffer content at stationarity has the same distribution as the random variable U_T , which is defined in terms of the discrete-time process $(U_t, \mathbf{V}_t)$ and the stopping time T , which have the following definition:

$$\begin{aligned} U_0 &= -L; \quad \mathbf{V}_0 = \overbrace{(0, 0, \dots, 0)}^L; \\ U_{t+1} &= U_t + P_t; \quad \mathbf{V}_{t+1} = \text{rotate}(\mathbf{V}_t) \vee (P_t, 0, \dots, 0), \end{aligned} \quad (1)$$

and

$$T = \min\{t : \mathbf{V}_t = \mathbf{1}\}. \quad (2)$$

A few words on the notation. The random vector $\mathbf{V}_t$ is a Boolean vector of length L (i.e., it consists only of zeros and ones). We denote the set of such vectors by $\mathbb{B}^L$. The function $\text{rotate}(\cdot)$ rotates a vector as follows:

$$\text{rotate}(v_1, v_2, \dots, v_L) = (v_2, \dots, v_L, v_1), \quad (3)$$

and the operator $\vee$ performs an elementwise logical or. Finally, P_t is the stationary ergodic error process, with $P_t = 0$ if an error occurs at time t , and $P_t = 1$ otherwise, where it must be understood that P_t runs backwards in time. The random variable U of the stationary buffer content hence corresponds in distribution to the random variable U_T .

It is easily verified that T is indeed a stopping time with respect to the process $(U_t, \mathbf{V}_t)$, i.e., it is adapted to the natural filtration of the process. The above process can be viewed as a variant of the classic ‘coupon-collector’ problem [12], which goes as follows. Let n objects be picked repeatedly according to a certain random process. Find the earliest time at which all n objects have been picked at least once. The difference between the problem at hand and the classic problem is that in the former, one attempts to get a specific object at any given time, in a cyclic manner, whereas in the latter one typically receives a random coupon at every attempt. In section 7 we show that in high-error environments, we obtain in the limit the familiar coupon-collector solution. Note that the same solution was obtained in [5], although without mentioning the relation with the coupon-collector problem.

With respect to the packet error process, we mainly consider two classes: (1) the classical independent and identically distributed (IID) error process, in which packet errors occur with a fixed error probability p ; and (2) finite state-space Markovian error processes, which have become the model of choice in many performance studies over wireless links [4, 3, 9, 20]. The large-deviations analysis can be extended to more general error processes as well, a fact we briefly elaborate on in section 4.

The Markovian error processes we consider in this paper consist of a discrete-time Markov chain (the background Markov chain) with a finite number of states M , which we require in the present work to be irreducible. Furthermore, each state has a distinct packet error probability. Let matrix $\mathbf{Q}$ denote the transition probability matrix of the background Markov chain, and let vectors $\mathbf{p}_0$ and $\mathbf{p}_1$ record respectively the error and success probabilities during each state. The vector π denotes the steady-state probability vector of the Markov chain. That is, π is the unique normalized vector which satisfies the equation $\pi = \pi \mathbf{Q}$. As is well-known, the transition matrix $\mathbf{Q}^{(r)}$ of a time-reversed Markov process is related to the transition matrix $\mathbf{Q}$ of the original process by the following formula:

$$\mathbf{Q}^{(r)} = \text{diag}(\pi)^{-1} \mathbf{Q}^T \text{diag}(\pi), \quad (4)$$




FIGURE 2. The alternative method to determine the resequencing buffer content, which accents the kinship with the classical coupon-collector problem. Erroneous transmissions are again indicated in gray. We travel backwards and count the number of successful transmissions until every group has at least one. Here, the dashed boxes represent successful transmissions that are not counted (that is, after time T). The buffer content is then found by subtracting the length of a roundtrip time. Hence, we find $16 - 6 = 10$ packets in the resequencing buffer.

where $\text{diag}(\mathbf{v})$ denotes the diagonal matrix constructed from vector $\mathbf{v}$ and $\mathbf{Q}^T$ is the transpose of $\mathbf{Q}$. In order to minimize notational clutter, we denote in the sequel by $\mathbf{Q}$ the transition matrix of the Markov process that travels *backwards* in time, as this will prove to be by far the most prevalent. We also introduce the matrices $\mathbf{P}_0$ and $\mathbf{P}_1$, defined as follows:

$$\mathbf{P}_0 = \text{diag}(\mathbf{p}_0)\mathbf{Q} \quad \text{and} \quad \mathbf{P}_1 = \text{diag}(\mathbf{p}_1)\mathbf{Q}. \quad (5)$$

3. Exact analysis of the receiver buffer content. In this section we present a method for the exact analysis of the buffer content distribution. That is, we want to find the distribution of U_t when $\mathbf{V}_t$ hits the state $(1, \dots, 1)$, (i.e. at stopping time T). This is related to so-called exit problems. While exit problems are about the distribution of the *time* until a certain set of states are hit, here we are interested in the value of an auxiliary process U_T at the hitting time.

We define a vector generating function $\mathbf{U}_{\mathbf{v}}(y, z)$ that keeps track of the time, of U_t , $\mathbf{V}_t$, and of the background state x_t of the error process:

$$[\mathbf{U}_{\mathbf{v}}(y, z)]_i = \sum_{t=0}^{\infty} y^t \mathbb{E}[z^{U_t} 1_{\{\mathbf{V}_t = \mathbf{v}, x_t = i\}}], \quad (6)$$

where notation 1_A denotes the indicator function, i.e., it is equal to 1 when event A is fulfilled and zero otherwise. Note that $\mathbf{U}_{\mathbf{v}}(y, z)$ is a 'transient' generating function, that is, it explicitly keeps track of the time t through the argument y .

After fairly straightforward manipulations in the transform domain, we can derive from the definition of the process (U_t, V_t) that the following set of equations holds:

$$U_0(y, z) = z^{-L}\pi + yU_0(y, z)P_0; \quad (7)$$

$$U_{1v}(y, z) = yzU_{v0}(y, z)P_1 + yU_{v1}(y, z)(P_0 + zP_1); \quad (8)$$

$$U_{0w}(y, z) = yU_{w0}(y, z)P_0; \quad (9)$$

$$U_1(y, z) = yzU_{10}(y, z)P_1, \quad (10)$$

for all $v \in \mathbb{B}^{L-1} \setminus \{1\}$ and for all $w \in \mathbb{B}^{L-1} \setminus \{0\}$.

Looking at the definition of the process (U_t, V_t) , we can deduce that the generating function $U(z)$ of the buffer content under steady-state conditions is equal to

$$U(z) = U_1(1, z)1. \quad (11)$$

In principle, the above derivations supply us with a system of 2^L equations from which the generating function $U(z)$ can be obtained. By way of example, we explicitly derive the buffer content distribution for the case $L = 2$. We have that

$$\begin{aligned} U_{00}(1, z) &= z^{-2}\pi(\mathbf{I} - P_0)^{-1}; \\ U_{10}(1, z) &= zU_{00}(1, z)P_1 + U_{01}(1, z)(P_0 + zP_1); \\ U_{01}(1, z) &= U_{10}(1, z)P_0; \\ U_{11}(1, z) &= zU_{10}(1, z)P_1. \end{aligned}$$

From this set of equations, we see that the buffer content distribution has the following generating function:

$$U(z) = \pi(\mathbf{I} - P_0)^{-1}P_1(\mathbf{I} - P_0^2 - zP_0P_1)^{-1}P_11. \quad (12)$$

Although for $L = 2$ the expression is simple enough, for larger L it simply becomes intractable to solve this system of 2^L equations. As we already stated in the introduction, this problem is severely affected by the curse of dimensionality. Note that, as in [5], we can circumvent this curse for IID error processes, as in that case we can perform some tricks to obtain an explicit expression for $U(z)$ for general L . The non-commutativity of matrices however prevents us from using these tricks in case of Markovian error processes.

As a numerical procedure, we find that this kind of analytical solution falls a bit short. In fact, for large L , it is very well possible that even a Monte-Carlo simulation of the stochastic process (U_t, V_t) can produce meaningful numerical results faster and with less memory usage than the solution sketched above, especially if that simulation program makes optimal use of importance sampling, or a similar variance reduction technique. In the next section we therefore show some computationally efficient analytical approximations.

4. Large-deviations analysis. Since the first steps were made in the works of Cramér in the 1930s, and especially after Donsker and Varadhan put it on a firm mathematical basis, large-deviations (LD) theory has always been a useful tool for practitioners in such diverse areas as risk management, statistical physics and tele-traffic engineering. Roughly speaking, large-deviations theory concerns the asymptotical behavior of stochastic processes, that is of situations that are large deviations from the mean. Typically, it can only tell something about the exponential decay rate of the probability of rare events, and as such it is not the most detailed of analytical tools. However, the decay rate of certain rare events is often the most

crucial performance measure, and furthermore, the power of large-deviations theory resides in the fact that it can be generalized rather easily, and that it can, with relative simplicity tell us something about systems for which other methods – including simulations – may fail.

In the following, we set out to compute the exponential decay rate for the buffer content distribution. That is, it is often observed (from simulations etc.) that for large k ,

$$\Pr[u > k] \approx Ce^{-k\delta}, \quad (13)$$

and we would like to compute the decay rate δ efficiently, (which is in practice of a far greater importance than C). Formally, δ is defined as the limit

$$\delta = - \lim_{k \rightarrow \infty} \frac{1}{k} \log \Pr[u > k], \quad (14)$$

provided that this limit exists.

We start out from the definition of the stochastic process $(U_t, \mathbf{V}_t)$. If U_T ends up with a value larger than k , it is because U_t reaches a value of k while the stopping time has not yet occurred, i.e., there exists a time instant t for which U_t becomes k and $\mathbf{V}_t \neq \mathbf{1}$. This leads to:

$$\Pr[U_T > k] = \sum_t \frac{1}{t} \Pr[U_{t-1} = k-1, \mathbf{V}_t \neq \mathbf{1}, P_{t-1} = 1]. \quad (15)$$

If the stopping time has not yet been reached, it means that no successful transmission has occurred in (at least) one of the groups. Let us define the modified process Y_t^s where parameter s indicates that group s is entirely erroneous:

$$Y_t^s = U_t 1_{\{P_t=1\}} \prod_{k \geq 0, kL+s < t} 1_{\{P_{kL+s}=0\}}. \quad (16)$$

We derive a large-deviations principle for process Y_t^s by means of the Gärtner-Ellis Theorem [10]. In order to do so, we define the scaled cumulant generating function $\Lambda_{s,t}(\theta)$:

$$\Lambda_{s,t}(\theta) \doteq \frac{1}{t} \log \mathbb{E}[e^{\theta Y_t^s}], \quad (17)$$

and the corresponding limit $t \rightarrow \infty$:

$$\Lambda_s(\theta) = \lim_{t \rightarrow \infty} \Lambda_{s,t}(\theta). \quad (18)$$

For Markovian error processes, $\Lambda_{s,t}(\theta)$ is equal to

$$\Lambda_{s,t}(\theta) = \log \pi \mathbf{P}(\theta)^{s-1} \mathbf{R}(\theta)^k \mathbf{P}(\theta)^r \mathbf{P}_1 e^\theta \mathbf{1}, \quad (19)$$

where k and r are uniquely defined by relations $t = s + kL + r + 1$ and $0 \leq r < L$. Also, we defined $\mathbf{P}(\theta) = \mathbf{P}_0 + e^\theta \mathbf{P}_1$, and $\mathbf{R}(\theta) = \mathbf{P}_0 \mathbf{P}(\theta)^{L-1}$. Note that when t goes to infinity, so does k , with $k \rightarrow L^{-1}t$ whereas s and r remain bounded. This enables us to invoke the Perron-Frobenius theory as follows:

$$\Lambda_s(\theta) = \lim_{k \rightarrow \infty} \frac{1}{kL} \log \mathbf{v}(\theta) \mathbf{R}(\theta)^k \mathbf{w}(\theta) \quad (20)$$

$$= \frac{1}{L} \log \rho(\mathbf{R}(\theta)), \quad (21)$$

where the notation $\rho(\cdot)$ denotes the spectral radius of a matrix. In the first step, we performed a change of variables in the limit, and assembled the asymptotically vanishing parts into the row vector $\mathbf{v}(\theta)$ and the column vector $\mathbf{w}(\theta)$. In the second step, we made use of a property that is proved in [10], §3.1, and that holds whenever

matrix $\mathbf{R}(\theta)$ is an irreducible matrix. Note that the spectral radius in that case is identical to the largest eigenvalue, and hence an intuitive explanation resides in the fact that in the limit, the contribution of the largest eigenvalue dominates all the other contributions. Also note that $\Lambda_s(\theta)$ is independent of s , and hence we drop the s when no ambiguity can arise and simply write $\Lambda(\theta)$.

The Gärtner-Ellis Theorem states that the rate function is equal to the Legendre-Fenchel transform of $\Lambda(\theta)$:

$$\Lambda^*(x) = \sup_{\theta} (\theta x - \Lambda(\theta)), \quad (22)$$

under certain regularity conditions on $\Lambda(\theta)$ (that is, it must be finite in a neighborhood of $\theta = 0$, essentially smooth and lower-semicontinuous). These conditions are fulfilled in this case, see eg. [14], p. 43.

The rate function $\Lambda^*(x)$, in turn, allows us to find the exponential decay of probabilities associated with process Y_t^s : for sufficiently 'nice sets' A (such as intervals) the following powerful result holds:

$$\Pr[t^{-1}Y_t^s \in A] = \exp(-t \inf_{x \in A} \Lambda^*(x) + o(t)), \quad (23)$$

where $o(t)$ denotes a function that decays faster than any linear function. Also note that the application of the Gärtner-Ellis Theorem is not restricted to the finite Markov error processes we are considering here. As long as the process Y_t^s admits a scaled cumulant function, for which the limit $\Lambda_s(\theta)$ exists and fulfills the above mentioned regularity conditions. Its computation can be performed either in closed form, numerically, or even by sampling (i.e., by a Monte-Carlo based technique).

Now, we are ready to derive the large-deviations result for the buffer content distribution proper.

Theorem 4.1. *The buffer content distribution satisfies*

$$\lim_{x \rightarrow \infty} \frac{1}{x} \log \Pr[U_T > x] = -\delta,$$

where $\delta = \inf_{x>0} \frac{\Lambda^*(x)}{x}$ or equivalently,

$$\delta = \sup\{\theta > 0 : \Lambda(\theta) < 0\}. \quad (24)$$

Proof. We start out by showing that the two definitions of the decay rate δ are equivalent. In fact, this is proved in many articles and books on large deviations, but for the sake of completeness, we repeat the proof here. Note that

$$\theta < \delta = \inf_{x>0} \frac{\Lambda^*(x)}{x} \text{ iff } \theta x - \Lambda^*(x) < 0 \text{ for } x > 0.$$

By the duality of the Legendre-Fenchel transform under convexity, we get

$$\Lambda(\theta) = \sup_{x>0} (\theta x - \Lambda^*(x)), \quad (25)$$

and hence, $\theta < \delta$, if and only if $\Lambda(\theta) < 0$, which leads to the desired

$$\delta = \sup\{\theta > 0 : \Lambda(\theta) < 0\}. \quad (26)$$

It is this last relation that proves to be far more useful. It amounts to finding a root of the cumulant generating function, without any need for Legendre-Fenchel transforms and so on.

Next, we derive a lower bound for the exponential decay rate. Observe that for any given s, t and x ,

$$\Pr[Y_t^s > x] \leq \Pr[U_T > x], \quad (27)$$

as the event $Y_t^s > x$ implies $U_T > x$. Hence, for any $c > 0$, and $x > 0$, and $s : 0 \leq s < L$,

$$\frac{1}{x} \log \Pr[U_T > x] \geq \frac{1}{c \lceil \frac{x}{c} \rceil} \log \Pr[\frac{Y_{\lceil \frac{x}{c} \rceil}^s}{\lceil \frac{x}{c} \rceil} > c], \quad (28)$$

where the notation $\lceil a \rceil$ denotes the smallest integer larger than a . Choose $t = \lceil \frac{x}{c} \rceil$, then we get

$$\begin{aligned} \liminf_{x \rightarrow \infty} \frac{1}{x} \log \Pr[U_T > x] \\ \geq \frac{1}{c} \liminf_{t \rightarrow \infty} \frac{1}{t} \log \Pr[\frac{Y_t^s}{t} > c] = -\frac{\Lambda^*(c)}{c}. \end{aligned} \quad (29)$$

Since c is chosen arbitrarily (as long as $c > 0$), we have indeed proved the lower bound

$$\liminf_{x \rightarrow \infty} \frac{1}{x} \log \Pr[U_T > x] \geq -\delta. \quad (30)$$

For the upper bound, we start from the inequality

$$\Pr[U_T > x] \leq \sum_{t \geq 1, 0 \leq s < L} \Pr[Y_t^s \geq x], \quad (31)$$

The reasoning behind this inequality is that if the event $\{U_T > x\}$ occurs, then there exists s, t for which $\{Y_t^s \geq x\}$. Hence, summing over all possible s and t , we get the desired inequality.

Using Chernoff's bound, we get:

$$\begin{aligned} \Pr[Y_t^s \geq x] &\leq \exp(-\theta x) \mathbb{E}[\exp(\theta Y_t^s)] \\ &\leq \exp(-\theta x + t \Lambda_t^s(\theta)), \end{aligned} \quad (32)$$

where in the second step we applied the definition of $\Lambda_t^s(\theta)$. Pick some $\theta > 0$ such that $\Lambda(\theta) < 0$ (we know that there exists such θ since $\mathbf{R}(0)$ is a substochastic matrix, and hence $\Lambda(0) < 0$ and $\Lambda(\theta)$ is differentiable in the neighborhood of $\theta = 0$). Choose $\epsilon > 0$ such that $\Lambda(\theta) + 2\epsilon\theta < 0$. Since $\Lambda_t^s(\theta) \rightarrow \Lambda(\theta)$, there exists t_0 such that for $t > t_0$

$$\Lambda_t^s(\theta) < \Lambda(\theta) + \epsilon\theta, \quad (33)$$

and hence

$$\Pr[U_T > x] \leq e^{-\theta x} \left(\sum_{1 \leq t \leq t_0, 0 \leq s < L} e^{t \Lambda_t^s(\theta)} + \sum_{t > t_0, 0 \leq s < L} e^{-\epsilon\theta t} \right). \quad (34)$$

The first sum is finite because each of the terms is finite, and the second sum is clearly finite. Hence

$$\limsup_{x \rightarrow \infty} \frac{1}{x} \log \Pr[U_T > x] \leq -\theta, \quad (35)$$

As we can choose θ as long as $\Lambda(\theta) < 0$ or equivalently $0 < \theta < \delta$, this concludes the proof. $\square$

Note that thanks to a careful choice of the processes $(U_t, \mathbf{V}_t)$ and Y_t , we were able to derive the decay rate δ along similar lines as for the standard $G/G/1$ queue, see for example the proofs in [14, 15, 18], which is somewhat remarkable as the workings of the resequencing buffer are very different. For example, it is hard to attribute a meaning to workload processes in resequencing buffers, which are the central tool in the LD derivation of standard queues.

For IID error processes (with error probability p), we are able to come up with an exact solution. Note that the cumulant generating function of Y_t^s is in this case equal to

$$\Lambda(\theta) = \frac{1}{L} \log(p(p + (1-p)e^\theta)^{L-1}), \quad (36)$$

so that

$$\delta = \log \frac{p^{-\frac{1}{L-1}} - p}{1-p}. \quad (37)$$

One of the appeals of large-deviations techniques is that the lengthiest derivations can often be summarized by a simple intuitive explanation. Here, this would be something along the lines of: in Selective Repeat resequencing buffers, large buffer contents occur only when there are many erroneous transmissions for one particular packet.

Note that the value of δ stays invariant under the reversion of the Markov chain. Indeed, let $\lambda(\theta)$ be an eigenvalue of $\mathbf{R}(\theta)$, and let $\mathbf{v}(\theta)$ be the corresponding left eigenvalue. When we construct the corresponding matrix expression $\mathbf{R}^{(r)}(\theta)$ for the reversed background Markov process $\mathbf{Q}^{(r)}$, we get

$$\begin{aligned} \mathbf{R}^{(r)}(\theta) &= \text{diag}(\boldsymbol{\pi})^{(-1)} \text{diag}(\mathbf{p}_0) \mathbf{Q}^T ((\text{diag}(\mathbf{p}_0) + e^\theta \text{diag}(\mathbf{p}_1)) \mathbf{Q}^T)^{L-1} \text{diag}(\boldsymbol{\pi}) \\ &= \text{diag}(\boldsymbol{\pi})^{(-1)} \text{diag}(\mathbf{p}_0) (\mathbf{P}(\theta)^{L-1})^T \mathbf{Q}^T \text{diag}(\boldsymbol{\pi}). \end{aligned} \quad (38)$$

We find that $\mathbf{w}(\theta) = \text{diag}(\boldsymbol{\pi})^{(-1)} \text{diag}(\mathbf{p}_0) \mathbf{v}(\theta)^T$ is a right eigenvector of $\mathbf{R}^{(r)}(\theta)$, corresponding to eigenvalue $\lambda(\theta)$. Indeed,

$$\mathbf{R}^{(r)}(\theta) \mathbf{w}(\theta) = \lambda(\theta) \text{diag}(\boldsymbol{\pi})^{(-1)} \text{diag}(\mathbf{p}_0) \mathbf{v}(\theta)^T. \quad (39)$$

As this holds for every eigenvalue $\lambda(\theta)$, this means that the spectra of $\mathbf{R}(\theta)$ and $\mathbf{R}^{(r)}(\theta)$ are identical, so that naturally the spectral radii are the same as well.

Summarizing the results of this section, in order to find the decay rate δ , we must find the smallest real root of an expression involving a spectral radius. By combining a good root-finding algorithm (such as the secant method or the Newton-Raphson method) with an efficient algorithm to compute the spectral radius (see e.g. [21]), this is a computationally feasible task, even for larger matrix dimensions.

5. A closer look at the large deviations limit. In the previous section, we looked at the *probability* that an exceptionally large buffer content occurs. In this section, we have a closer look at how this rare event exactly occurs. Indeed, we know that one group will stay erroneous during a large time, but what is the expected number of successful transmissions in the other groups? Is it larger or smaller than during ‘typical’ runs of the process? Is it the same for all groups? These questions are not merely of theoretical interest, but they are also important for the design of efficient simulation algorithms.

For an IID error processes, the error rate during a build-up to the rare event is of course the same for every group except the erroneous one. In our finite Markov-chain setup, the situation is not as simple. An equal error rate in each of the groups is not assured. We introduce the $L-1$ -dimensional functions $\Lambda(\boldsymbol{\theta})$ and $\Lambda^*(\mathbf{x})$, which are respectively the limiting cumulant generating function and the rate function. They are defined as follows:

$$\Lambda(\boldsymbol{\theta}) = \log \rho(\mathbf{P}_0 \mathbf{P}(\boldsymbol{\theta}_1) \cdots \mathbf{P}(\boldsymbol{\theta}_{L-1})) \quad \text{and} \quad \Lambda^*(\mathbf{x}) = \sup_{\boldsymbol{\theta}} (\boldsymbol{\theta} \cdot \mathbf{x} - \Lambda(\boldsymbol{\theta})). \quad (40)$$

Intuitively, the rate function $\Lambda^*(x_1, x_2, \dots)$ gives the decay rate when the success rate during the first group is x_1 , during the second, x_2 , and so on. We thus bring the large-deviations principle to a more detailed space.

We bring the definition (24) of δ also into this finer space as follows:

$$\delta = \sup\{\theta > 0 : \Lambda(\theta, \dots, \theta) < 0\}. \quad (41)$$

As eigenvalues are differentiable in a parameter whenever the original matrix is differentiable in that parameter, we have that $\Lambda(\theta)$ is differentiable in the relevant domain. Hence, the Legendre-Fenchel transform of $\Lambda(\theta)$ reduces to the simpler Legendre transform, and in particular,

$$\Lambda^*(\mathbf{x}) = (\delta, \dots, \delta) \cdot \mathbf{x} - \Lambda(\delta, \dots, \delta), \quad \text{where } \mathbf{x} = \nabla \Lambda(\theta)|_{\theta=(\delta, \dots, \delta)}. \quad (42)$$

It is exactly this vector $\mathbf{x}$ in which we are interested. With the help of some matrix manipulations we get simple formulas for $\mathbf{x}$. Indeed, its i th component is equal to

$$\begin{aligned} x_i &= \frac{\partial \Lambda}{\partial \theta_i} |_{\theta=(\delta, \dots, \delta)} \\ &= \frac{1}{\rho(\mathbf{R}(\delta))} \frac{\partial}{\partial \theta_i} \rho(\mathbf{P}_0 \mathbf{P}(\theta_1) \dots \mathbf{P}(\theta_{L-1})) |_{\theta=(\delta, \dots, \delta)}. \end{aligned} \quad (43)$$

Due to the definition of δ , the fraction is equal to 1, whereas the remaining factor can be rewritten by the formula for the derivative of an eigenvalue. Indeed, let ℓ and $\mathbf{r}$ denote the left and right eigenvectors corresponding to the Perron-Frobenius eigenvalue of $\mathbf{R}(\delta)$, then we have that

$$\begin{aligned} x_i &= \ell \frac{\partial}{\partial \theta_i} (\mathbf{P}_0 \mathbf{P}(\theta_1) \dots \mathbf{P}(\theta_{L-1})) \mathbf{r} \\ &= \ell \mathbf{P}_0 \mathbf{P}(\delta)^{i-1} \mathbf{P}_1 \mathbf{P}(\delta)^{L-i-1} \mathbf{r}. \end{aligned} \quad (44)$$

Hence, we found that in sample paths leading to large buffer contents, the success rate in each of the groups is not necessarily equal but given by equation (44). After some trivial manipulations, we find that for IID error processes, the success rate is the same in each of the groups, and equal to $1 - p^{\frac{L}{L-1}}$, which is (slightly) higher than during a typical run.

The relationship between large-deviations results and exponentially tilted importance sampling is well-known, see e.g. [2]. From the fact that the optimal decay rate δ is reached in $\Lambda(\delta, \dots, \delta)$ (with all arguments the same), we see that the optimal change of measure tilts the Markov chain equally with parameter δ in all of the groups, except for the first group, where it is tilted with parameter $-\infty$. We have put this knowledge to use to improve the efficiency of our simulation program.

6. Asymptotic analysis of the packet delay. We now tackle the asymptotics of the packet delay. The heuristic reasoning is as follows. An exceptionally long delay of a packet is only possible when there is a long burst of errors in one of the groups. Let D be the random variable denoting the delay of an arbitrary packet. Then we postulate the following on the decay rate ν of the packet delay:

Theorem 6.1. *The packet delay distribution has an exponential asymptotic with decay rate ν :*

$$\lim_{k \rightarrow \infty} \frac{1}{k} \log \Pr[D > k] = -\nu,$$

where

$$\nu = -\frac{1}{L} \log \rho(Q^{L-1}P_0). \quad (45)$$

Proof. We find lower and upper bounds of the probability for the event $\{D > k\}$, and show that they have the same exponential asymptotic.

We start with an upper bound. If we have a delay larger than k , then necessarily, there is a group with minimally $\lceil (k-1)/L \rceil$ consecutive errors. Let i be the channel state of the slot during which series of errors starts. For the upper bound we are interested in the channel state that maximizes the probability. Thus we have an upper bound:

$$\begin{aligned} \Pr[D > k] &\leq \sup_i \Pr[\# \text{ errors in a group} = \lceil (k+1)/L \rceil] \\ &= \sup_i \mathbf{e}_i(Q^{L-1}P_0)^{\lceil (k+1)/L \rceil} \mathbf{1}, \end{aligned} \quad (46)$$

where $\mathbf{e}_i$ is a row vector for which the i th component is 1, and the other components are zero. From the Perron-Frobenius theorem it is hence easily seen that

$$\lim_{k \rightarrow \infty} \frac{1}{k} \log \Pr[D > k] \leq -\nu. \quad (47)$$

Now we look at the lower bound. We construct a superset E of $\{D > k\}$ in the event space, so that when E holds, then necessarily also $\{D > k\}$. The probability of E is then smaller than $\Pr[D > k]$, giving us the desired lower bound. Informally speaking, we need two ingredients so that a sample path produces a packet delay larger than k : (1) there must be a packet that is erroneously transmitted for a large number of times, say n , where $n \geq \lceil (k+1)/L \rceil + 1$, and (2) there must be an arrival in the reordering buffer of a packet that is younger than said packet. This means we have a situation as in figure 3. But since we track the delay of an arbitrary packet, we condition on the event that the arrival indeed occurs. Also note that for the lower bound, we look for the combination of channel state i and 'gap' j that minimizes the probability. Hence we have:

$$\begin{aligned} \Pr[D > k] &\geq \inf_{i,j,m} \Pr[n \text{ errors} | \text{packet arrives with certain } i, j, m] \\ &= \inf_{i,j,m} \mathbf{e}_i P_1 Q^{j-1} P_0 (Q^{L-j-1} P_0 Q^{j-1} P_0)^{m-1} Q^{L-j-1} P_1 Q^{j-1} P_0 \times \\ &\quad \times (Q^{L-1} P_0)^{n-1} \mathbf{1} \times \\ &\quad \times (\mathbf{e}_i P_1 Q^{j-1} P_0 (Q^{L-j-1} P_0 Q^{j-1} P_0)^{m-1} Q^{L-j-1} P_1 \mathbf{1})^{-1}. \end{aligned} \quad (48)$$

For large k , the contribution of the matrix power on the middle line dominates and hence we can apply the Perron-Frobenius theorem as before. Even if m is large, it cannot dominate as the terms in numerator and denominator cancel each other. Hence we have the lower bound

$$\lim_{k \rightarrow \infty} \frac{1}{k} \log \Pr[D > k] \geq -\nu, \quad (49)$$

and this concludes the proof. $\square$

7. Approximations for high error probabilities. When the error probabilities approach one, the time between two error-free transmissions will be sufficiently high, so that (1) the groups of successive error-free transmissions are almost independent, and (2) the Markov chain of the channel has reached stationarity. Under such conditions, we can directly apply the results of the coupon-collector problem.




FIGURE 3. Necessary event for the occurrence of a large packet delay. As before, white blocks denote correct transmissions, gray blocks denote incorrect ones and dashed blocks denote transmissions which may be either. The variables j , m and n are as defined in the text.

The value of interest in the original problem is the time τ (which in our problem translates to the number of successful transmissions) until we have collected all coupons. Under the above stated conditions, the buffer content during regime U_T can be approximated as $U_T = \tau - L$. We mention the mean, variance and tail probabilities (the derivations can be found in most textbooks on combinatorics, e.g. in [12]):

$$\begin{aligned} E[U_T] &= L(H_L - 1) \approx L \log L + \gamma L + \frac{1}{2} \\ \text{Var}[U_T] &= L^2 H_{L-1}^{(2)} - L H_{L-1} \approx \frac{\pi^2}{6} L^2 \\ \Pr[U_T > k] &\approx e^{-k \frac{L}{L-1}}. \end{aligned} \quad (50)$$

In these equations, H_L denotes the L th harmonic number, and $H_L^{(2)}$ the L th second order harmonic number. By γ we denote the Euler-Mascheroni constant: $\gamma \approx 0.5772157 \dots$.

Note that the mathematical elegance of this high-error approximation surpasses its practical usability, as consensus seems to be that the usage of retransmission protocols over channels with an extremely high packet error rate leads to poor performance.

8. Numerical results and discussion. In this section, we show the validity of our large-deviations technique by means of some specific numerical examples. First, we consider some cases where the error process is IID. In Figure 4, the decay rate δ of the buffer content distribution is plotted against the packet error probability p , for different values of the feedback delay L . We notice that the decay rate is higher




FIGURE 4. Plot showing the influence of the error probability p on the decay rate δ of the buffer content distribution, for different values of the feedback delay L .




FIGURE 5. Plot showing the exact distribution of the buffer content, obtained by simulations (with importance sampling; full lines) and also the tail approximations obtained by large deviations (dashed lines), for $p = 0.3$ and different values of the feedback delay L .

for small error probabilities, which is consistent with the intuition that in that case, higher buffer contents are very unlikely. Also as intuitively expected, the decay rate decreases with increasing values of the feedback delay. For $p \rightarrow 1$, the decay rate approaches the value of $\log \frac{L}{L-1}$ that is predicted by the high error rate analysis.

In Figure 5, we show that our model indeed predicts the decay rates right. We repeat that large deviations in general predict only the decay rate and not the offset of the tail, and therefore, we have chosen an offset of zero (corresponding to $C = 1$ in the form $C \exp(-\delta n)$). We obtain accurate simulation results with the help of importance sampling.

In the next pair of plots, we show some results in case of Markovian error processes. We take a three-state error model that has been derived in [20] by means of the Hidden Markov estimation techniques. They found for a fading margin $F = 20\text{dB}$ and a Doppler frequency normalized to the slot length of $f_D T = 0.01$,




FIGURE 6. Plot showing the influence of error correlation on the decay rate δ of the buffer content distribution.

the following matrices:

$$\begin{aligned}
 P_0 &= \begin{pmatrix} 0 & 0.002708 & 0.000213 \\ 0 & 0.646523 & 0 \\ 0 & 0 & 0.878338 \end{pmatrix}, \\
 P_1 &= \begin{pmatrix} 0.997079 & 0 & 0 \\ 0.353477 & 0 & 0 \\ 0.121662 & 0 & 0 \end{pmatrix}.
 \end{aligned} \tag{51}$$

In Figure 6, we show the influence of the correlation on the decay rate. That is, we compare the decay rate of the Markovian channel with the decay rate of the corresponding IID channel (i.e. with the same time-average error probability, which in this case is given by $p = 0.009324$). We see that especially for smaller feedback delays, there is a considerable difference between the correlated and uncorrelated error process. In Figure 7, we compare the obtained asymptotics with simulations.

Finally, in figure 8, we show the decay rates of buffer content and the packet delay, for both the Markov process and the corresponding IID process. A fairly remarkable fact is that although for smaller roundtrip times there are important differences in decay rates, for very large decay rates the four curves converge to the same asymptote.

9. Conclusions. We have analyzed the performance of the resequencing buffer for the selective-repeat protocol over correlated channels. We found an exact solution for the buffer content distribution over Markovian error channels, as well as computationally efficient large-deviations results. We touched upon the similarity with the coupon-collector problem, to which it reduces in the limit as packet error probabilities increase. Comparisons with simulations confirm the validity of our approach.

Acknowledgments. The authors would like to thank the anonymous referees for their valuable comments, especially with respect to the proof of Theorem 1.




FIGURE 7. Plot showing the exact distribution of the buffer content (obtained by simulations) versus the LD approximations for the 3 state Markov model (HMM) and the corresponding IID error process. The feedback delay L equals 5.




FIGURE 8. Plot showing the decay rates of the packet delay and the buffer content, each for IID and Markovian error processes.

REFERENCES

- [1] M. Anagnostou and E. Protonarios, *Performance analysis of the selective repeat protocol*, IEEE Transactions on Communications, 34 (1986), 127–135.
- [2] S. Asmussen and P. W. Glynn, "Stochastic Simulation: Algorithms and Analysis," volume 57 of "Stochastic Modelling and Applied Probability," Springer, New York, 2007.
- [3] L. Badia, *On the impact of correlated arrivals and errors on ARQ delay terms*, IEEE Transactions on Communications, 57 (2009), 334–338.
- [4] L. Badia, M. Rossi and M. Zorzi, *SR ARQ packet delay statistics on Markov channels in the presence of variable arrival rate*, IEEE Transactions on Wireless Communications, 5 (2006), 1639–1644.
- [5] H. Bruneel, J. De Vriendt and C. Ysebaert, *Receiver buffer behavior for the selective-repeat ARQ protocol*, Comput. Netw. ISDN Syst., 19 (1990), 129–142.
- [6] S. R. Chakravorthy and S. Chukova, *A finite capacity resequencing model with Markovian arrivals*, Asia-Pac. J. Oper. Res., 22 (2005), 409–443.

- [7] S. De Turck and Koen Wittevrongel, *Delay analysis of Go-Back-N ARQ for correlated error channels*, pages 245–267. River Publishers, Aalborg, 2009.
- [8] S. De Vuyst, K. Tworus, S. Wittevrongel and H. Bruneel, *Analysis of stop-and-wait ARQ for a wireless channel*, 4OR, **7** (2009), 61–78.
- [9] S. De Vuyst, S. Wittevrongel and H. Bruneel, *Queueing delay of stop-and-wait ARQ over a wireless Markovian channel*, Mobility Management and Quality-of-Service for Heterogeneous Networks, (2009), 1–31.
- [10] A. Dembo and O. Zeitouni, “Large Deviations Techniques and Applications,” Second edition, volume 38 of “Applications of Mathematics” (New York), Springer-Verlag, New York, second edition, 1998.
- [11] R. Fantacci, *Queueing analysis of the selective repeat automatic repeat request protocol wireless packet networks*, IEEE Trans. Veh. Technol, **45** (1996), 258–264.
- [12] P. Flajolet and R. Sedgewick, “Analytic Combinatorics,” Cambridge University Press, Cambridge, 2009.
- [13] S. Fujii and Y. Hayashida, *Delay performance of a selective-repeat ARQ scheme with a Markovian error channel*, Electronics and communications in Japan Part I - Communications, **81** (1998), 31–41.
- [14] A. Ganesh, N. O’Connell and D. Wischik, “Big Queues,” Lecture Notes in Mathematics, 1838, Springer-Verlag, Berlin, 2004.
- [15] P. W. Glynn and W. Whitt, *Logarithmic asymptotics for steady-state tail probabilities in a single-server queue*, Studies in Applied Probability, J. Appl. Probab., **31A** (1994), 131–156.
- [16] J. G. Kim and M. Krunk, *Delay analysis of selective repeat ARQ for a Markovian source over a wireless channel*, IEEE Trans. Veh. Technol, **49** (2000), 1968–1981.
- [17] A. Konheim, *A queueing analysis of two ARQ protocols*, IEEE Trans. Commun., **10** (1992), 1004–1014.
- [18] E. Morozov, “Communications Systems: Rare Events Simulation and Effective Bandwidths,” Public University of Navarre, 2004.
- [19] Z. Rosberg and N. Shacham, *Resequencing delay and buffer occupancy under the selective-repeat arq*, IEEE Trans. Inform. Theory, **35** (1989).
- [20] W. Turin and M. Zorzi, *Performance analysis of delay-constrained communications over slow Rayleigh fading channels*, IEEE Trans. Wireless Comm., **1** (2002), 801–807.
- [21] R. J. Wood and M. J. O’Neill, *Finding the spectral radius of a large sparse non-negative matrix*, ANZIAM J., **48** (2007), C330–C345 (2008).
- [22] Y. Xia and D. Tse, *Analysis on packet resequencing for reliable network protocols*, Perform. Eval., **61** (2005), 299–328.

Received September 2009; 1st revision January 2010; final revision April 2010.

E-mail address: kdeturck@telin.ugent.be

E-mail address: sw@telin.ugent.be