# Randomly Select and Forward: Erasure Probability Analysis of a Probabilistic Relay Channel Model

Alberto Tarable*, Ioannis Chatzigeorgiou† and Ian J. Wassell†
*DELEN, CERCOM-Politecnico di Torino
Email: alberto.tarable@polito.it
†Computer Laboratory, University of Cambridge
Email: {ioannis.chatzigeorgiou, ian.wassell}@cl.cam.ac.uk

***Abstract*— This paper proposes a semi-analytical framework for estimating the erasure probability in single-hop multi-relay networks. Specifically, we consider a system, in which relays do not decode the information but simply forward coded packets that have been previously received from the source. This allows for uncoordinated, low-complexity processing at the relays. We present a detailed analysis of the proposed network model, which represents an instance of a new coding scheme that we refer to as distributed fountain coding.**

## I. Introduction

Thanks to their versatility, fountain codes [1], [2] are a natural coding solution for many network topologies, including broadcast, multicast and many-to-one communications. This paper describes our first effort to explore the application of fountain codes to single-hop multi-relay networks. More specifically, our system consists of a source that conveys coded information to a destination, assisted by a bank of half-duplex relays. Throughout the paper it is assumed that there is no direct link between the source and the destination.

The application of fountain coding principles to various topologies of relay networks has also been studied in [3], [4] and [5]. In particular, each relay of the scheme described in [3] acts like a receiver, until it decodes the source information; then, it switches to a transmitter and broadcasts coded packets, which are generated by the same fountain code used at the source. Note that each receiving relay tries to decode the source information by collecting coded packets from both the source and the transmitting relays. In [4], a simple single-relay scheme is described and information-theoretic tools are used to assess its performance. In [5], a two-relay scheme is presented. In odd time steps, one relay transmits while the other receives; in even time steps, their roles are reversed. At any time step, the source and the transmitting relay forward coded packets to both the destination and the receiving relay. It is important to note that a packet transmitted by a relay is generated by modulo-2 adding (i.e., XORing) a random number of coded packets, which have been successfully received in previous time steps and stored in the memory of the relay. Thus, unlike the previous schemes, relays do *not* decode the source information; they only perform what can be called "linear network coding". Consequently, the computational complexity at the relays is considerably reduced.

In this paper, we present a generalization of the scheme described in [5]. We consider not only two but an arbitrary number of relays. Furthermore, we do not use a pre-agreed protocol that determines which relays transmit and which relays receive; instead, each relay flips a possibly biased coin at each time step and, according to the result, it either transmits a coded packet to the destination or receives a coded packet from the source. The main advantage of this probabilistic approach is that no centralized control is required, since each relay operates independently of the others.

Nevertheless, in order to simplify our analysis, we have assumed that a transmitting relay randomly selects and forwards one of the coded packets that have been stored in its memory. This is a particular case of a more general scheme, in which relays are allowed to XOR stored coded packets. In other words, in this paper we study an instance of what could be called a *distributed fountain code*.

The paper is organized as follows: Section II introduces the system model. In Section III and Section IV, we give a semi-analytical method, based on Markov chains, to estimate the main performance parameter of our proposed scheme, namely the erasure probability. In Section V, we compare theoretical to simulation results and discuss the impact of the scheme parameters on the overall erasure probability. Finally, the main conclusions of this work are summarized in Section VI.

## II. System Model

The basic system setup is depicted in Fig. 1. We consider a wireless network that consists of a source node (S), a total of $M$ relay nodes and a destination node (D); all nodes are equipped with a single antenna.

The source uses a binary random linear fountain code to encode $K$ packets, denoted as $u_1, u_2, \ldots, u_K$. At each time step, labeled by $n$, the encoder generates a coded packet $c_n$, which is transmitted to the relays and the destination. It is important to underline that the use of a fountain code at the source is not a prerequisite for our analysis; the choice of the transmission scheme does not affect our framework.

At time step $n$, some relays listen to the source with probability $p_{\mathrm{L}}$, while the remaining relays transmit to the destination. The relays that listen, perform a simple cyclic redundancy check (CRC) on the received packet $c_n$. A listening relay that has successfully recovered $c_n$, stores it in a buffer that can hold up to $\nu$ packets. If CRC decoding was

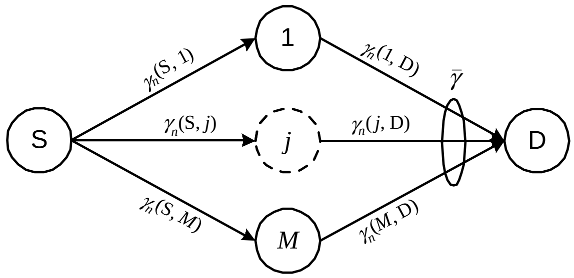


Fig. 1. Wireless network configuration with a single source, $M$ relays and a single destination.

unsuccessful, $c_n$ is dropped. Each relay that is in transmitting mode, randomly selects a coded packet from its memory and forwards it to the destination. We assume that relays use properly separated directional antennas, so that inter-relay interference is prevented at the relays but not at the destination.

The destination performs successive interference cancelation (SIC) to recover the received coded packets that were sent by the source and the relays over block fading channels. When $K$ coded packets have been transmitted by the source in an equal number of time steps, the destination starts to check whether the received packets can be decoded. If the destination successfully decodes the source packets after $N \geq K$ time steps, it signals the source to cease transmission.

The quality of a channel in our system model is characterized by the corresponding average receive signal-to-noise ratio (SNR). In particular, we assume that the fading coefficient of a channel between nodes $i$ and $j$ at time step $n$ is described by $h_n(i,j)$, which is modeled as a zero-mean, circularly symmetric complex Gaussian random variable with variance $\sigma^2(i,j)=1$. Each channel is also impaired by additive white Gaussian noise of variance $N_0$. Consequently, the instantaneous receive SNR at node $j$ is given by

$$\gamma_n(i,j) = |h_n(i,j)|^2 \frac{P_{\mathrm{T}}(i)}{N_0}, \tag{1}$$

where $P_{\mathrm{T}}(i)$ is the transmit power of node $i$. Considering that the average SNR is defined as $\overline{\gamma}(i,j) \triangleq \mathbb{E}\big[\gamma_n(i,j)\big]$, we obtain

$$\overline{\gamma}(i,j) = \mathbb{E}\Big[\big|h_n(i,j)\big|^2\Big] \frac{P_{\mathrm{T}}(i)}{N_0} = \frac{P_{\mathrm{T}}(i)}{N_0}, \tag{2}$$

where $\mathbb{E}\,[.]$ denotes the expectation operator.

In order to simplify the analysis, we assume throughout this paper that:

- The channels between the source S and the relays are perfect, that is $\overline{\gamma}(\mathrm{S},j) \to \infty$, for $j=1,2,\ldots,M$.
- All nodes have equal transmit power $P_{\mathrm{T}}(i) = P_{\mathrm{T}}$ and, consequently, the relay-to-destination channels are statistically similar, that is $\overline{\gamma}(j,\mathrm{D})=\overline{\gamma}$.
- The destination cannot obtain information directly from the source, that is $\overline{\gamma}(\mathrm{S},\mathrm{D})=0$.

Note, however, that our proposed framework can be extended in a straightforward manner to relay networks, in which the average SNR between the source and the relays or the source and the destination take finite, non-zero values.

## III. System Analysis using Markov Chains

In this section, we describe the transmission process of a packet from the source to the destination using a Markov chain; furthermore, we exploit the properties of Markov chains to compute the erasure probability of the packet.

### A. Definitions for a system using a single relay of memory $\nu$

We first consider the simple case when a single relay of memory size $\nu$ is used. We can represent the life cycle of a coded packet in the network, from transmission to successful decoding or erasure, using an absorbing Markov chain as shown in Fig. 2. In particular, $S_{\mathrm{I}}$ is the *initial state* of the chain, which represents the instant that the source transmits the packet, and $S_{\mathrm{S}}$ is the *state of success*, which signifies successful decoding of the packet at the destination. Every other Markov state $s$ can be described by a nonnegative integer $\xi$, which identifies the position that the packet of interest occupies in the memory of the relay; that is, $s=\xi$, where $\xi=0,1,\ldots\,\nu$. The packet is initially stored in the first memory position $(s=1)$ but it is shifted to the following positions as soon as new packets are received. If the zero state $s=0$ is reached, an erasure occurs. Note that the zero state and the state of success are the absorbing states of the Markov chain, while all other states are known as transient states.

The probability of transition between two states depends on the probability of the relay being in listening or transmitting mode, that is $p_{\mathrm{L}}$ or $1-p_{\mathrm{L}}$, respectively. Furthermore, if the relay is in transmitting mode, the transition probability should also be weighted by the probability that the destination is unsuccessful in recovering the packet of interest. We denote the latter probability as $f(k_1,k_2)$, where $k_1$ is the number of relays that have the packet stored in their memory, whilst $k_2$ is the number of those relays among the $k_1$ relays that are in transmitting mode. In the remainder of this section, we refer to a transition probability that should have been weighted by the probability of unsuccessful decoding, yet it has not been multiplied by $f(k_1,k_2)$, as *unweighted*.

Before proceeding to the general case of a network of $M$ relays, we first define the following matrices that will help us construct the transition matrix of the complete Markov chain. Let $\mathbf{V}^{(1)}=(V^{(1)}_{i,j})$ be a $(\nu{+}1)\times(\nu{+}1)$ matrix, specific for the single-relay scenario; its $(i,j)$ entry corresponds to the unweighted transition probability from state $s_i$ to state $s_j$, where $(s_1,s_2,\ldots,s_\nu,s_{\nu+1}) = (1,2,\ldots,\nu,0)$. Based on Fig. 2, we obtain

$$V^{(1)}_{i,j} = \begin{cases} p_{\mathrm{L}}, & \text{if } i=j-1 \\ 1-p_{\mathrm{L}}, & \text{if } i=j<\nu+1 \\ 1, & \text{if } i=j=\nu+1 \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

Similarly, let $\mathbf{v}^{(1)}=(v^{(1)}_j)$ be a row vector of length $\nu+1$, specific for the case when $M=1$. The $j$-th entry of the vector, which conveys the unweighted probability of transition from the initial state $S_{\mathrm{I}}$ to state $s_j$, can only take one of the

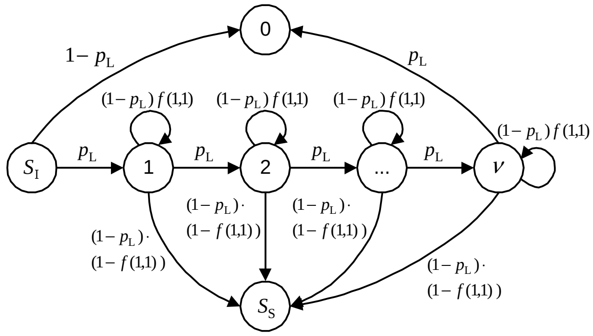


Fig. 2. Markov chain for a single-relay network ($M=1$).

following values

$$v_j^{(1)} = \begin{cases} p_{\mathrm{L}}, & \text{if } j=1 \\ 1-p_{\mathrm{L}}, & \text{if } j=\nu+1 \\ 0, & \text{otherwise.} \end{cases} \tag{4}$$

### B. Generalization for networks of $M$ relays

We now consider the general case of a network that uses $M$ relays of memory size $\nu$ each. The corresponding Markov chain consists of the initial state $S_{\mathrm{I}}$ and the state of success $S_{\mathrm{S}}$ as well as $\wp=(\nu+1)^M$ states of the form

$$s=(\xi_1,\xi_2,\ldots,\xi_M). \tag{5}$$

Here $\xi_k,\ k=1,\ldots,M$, denotes the position in the memory of the $k$-th relay in which the packet is stored. As previously mentioned, $\xi_k$ takes values between 0 and $\nu$. An erasure occurs when the packet is flushed from the memory of all relays and the Markov chain is absorbed by state $s=(0,\ldots,0)$, which we refer to as the *all-zero state*.

The unweighted transition probabilities in the Markov process of a $M$-relay network can also be expressed in the form of two matrices, namely $\mathbf{V}^{(M)}$ and $\mathbf{v}^{(M)}$, which have dimensions $\wp\times\wp$ and $1\times\wp$ respectively. It can be shown that $\mathbf{V}^{(M)}$ and $\mathbf{v}^{(M)}$ are directly related to matrices $\mathbf{v}^{(1)}$ and $\mathbf{V}^{(1)}$, presented in (3) and (4), as follows

$$\begin{aligned} \mathbf{V}^{(M)} &= \overbrace{\mathbf{V}^{(1)}\otimes\mathbf{V}^{(1)}\otimes\ldots\otimes\mathbf{V}^{(1)}}^{M \text{ times}} \\ &= \left(\mathbf{V}^{(1)}\right)^{[M]} \end{aligned} \tag{6}$$

and

$$\mathbf{v}^{(M)} = \left(\mathbf{v}^{(1)}\right)^{[M]} \tag{7}$$

where $\otimes$ and $[\cdot]$ denote the Kronecker product and Kronecker exponentiation, respectively.

We now introduce two new matrices, denoted as $\mathbf{b}=(b_j)$ and $\mathbf{B}=(B_{i,j})$, which are the weighted versions of $\mathbf{v}^{(M)}$ and $\mathbf{V}^{(M)}$. However, weighting has no effect on $\mathbf{v}^{(M)}$, since the transition probabilities from the initial state to any other state only depend on $p_{\mathrm{L}}$, hence

$$\mathbf{b}=\mathbf{v}^{(M)} \tag{8}$$

always. On the other hand, weighting will affect the value of some elements of $\mathbf{V}^{(M)}$, depending on the start state and the end state of the corresponding transitions. More specifically, the $(i,j)$ entry of $\mathbf{B}$ will assume the following value

$$B_{i,j} = \begin{cases} V_{i,j}^{(M)}, & \text{if } w_{\mathrm{H}}(s_i)=d_{\mathrm{H}}(s_i,s_j) \\ V_{i,j}^{(M)} f\big(w_{\mathrm{H}}(s_i), w_{\mathrm{H}}(s_i)-d_{\mathrm{H}}(s_i,s_j)\big), & \\ & \text{if } w_{\mathrm{H}}(s_i)>d_{\mathrm{H}}(s_i,s_j), \end{cases} \tag{9}$$

where $w_{\mathrm{H}}(s)$ denotes the Hamming weight of a state $s$ and $d_{\mathrm{H}}(s,s')$ represents the Hamming distance between states $s$ and $s'$. The single-relay case $(M=1)$ can be easily confirmed by examining Fig. 2.

### C. Derivation of the erasure probability

Having obtained matrices $\mathbf{b}$ and $\mathbf{B}$, we can construct the complete $(\wp+2)\times(\wp+2)$ transition matrix $\mathbf{T}$ as follows

$$\mathbf{T} = \left(\begin{array}{c|c|c} \mathbf{0}_{(\wp+1)\times 1} & \begin{array}{c}\mathbf{b}\\ \mathbf{B}\end{array} & \begin{array}{c}\mathbf{r}_{\mathrm{S}}\\ \hline 0\end{array} \\ \hline \multicolumn{2}{c|}{\mathbf{0}_{1\times(\wp+1)}} & 1 \end{array}\right). \tag{10}$$

Note that $\mathbf{r}_{\mathrm{S}}$ is a column vector of length $\wp$ that provides the probabilities of transition from all transient states to the state of success, $S_{\mathrm{S}}$, and can be easily evaluated, considering that the elements of each row in $\mathbf{T}$ should sum to one. Furthermore, each element of the first column of $\mathbf{T}$ corresponds to the probability of transition from a particular state to the initial state $S_{\mathrm{I}}$, which is always zero (see Fig. 2).

If we regroup the elements of $\mathbf{T}$, we can express the transition matrix in canonical form [6]

$$\mathbf{T} = \left(\begin{array}{c|c} \mathbf{Q} & \mathbf{r_0} \quad \mathbf{r}_{\mathrm{S}} \\ \hline \mathbf{0}_{2\times\wp} & \mathbf{I}_2 \end{array}\right), \tag{11}$$

where $\mathbf{Q}$ is a $\wp\times\wp$ matrix that contains the transition probabilities between transient states, whilst the entries of the length-$\wp$ column vector $\mathbf{r_0}$ correspond to all probabilities of transition from the transient states to the all-zero state.

The probability that the Markov chain will be absorbed by the all-zero state, starting from any of the $\wp$ transient states, is expressed in the form of column vector $\mathbf{t}=[t_1,t_2,\ldots,t_\wp]$, which can be obtained as follows [6]

$$\mathbf{t}=\left(\mathbf{I}_\wp-\mathbf{Q}\right)^{-1}\mathbf{r_0}. \tag{12}$$

The probability that a transmitted packet will be erased, given that the Markov process originated from the initial state, is conveyed by the first element of $\mathbf{t}$, that is

$$P_{\mathrm{e}}=t_1. \tag{13}$$

## IV. Computation of $f(k_1,k_2)$

As mentioned in Subsection III-A, $f(k_1,k_2)$ is the probability that the destination does *not* successfully recover the packet of interest, given that there are $k_1$ relays possessing the packet in their memories and that $k_2$ out of them are in transmitting mode. However, before deriving an expression for $f(k_1, k_2)$, we first evaluate the SIC performance of the system, which can be quantified by the probability that the destination has successfully recovered some of the packets that were received at a particular time step.

*Proposition 1:* Let $T$ be the number of packets that were transmitted to the destination by an equal number of relays at a given time step. The probability that the destination will recover $\ell$ packets out of the $T$ received packets is

$$q_\ell^T = \beta_\ell - \beta_{\ell+1}, \qquad \text{for} \;\; \ell = 0, 1, \ldots, T \tag{14}$$

where $\beta_{T+1} = 0$ and

$$\beta_\ell = \frac{T!}{(T-\ell)!} \, 2^{-\frac{\ell(2T-\ell-1)}{2}} \, e^{-\frac{2^\ell - 1}{\overline{\gamma}}}. \tag{15}$$

*Proof:* If $\gamma_1, \gamma_2, \ldots, \gamma_T$ are the instantaneous SNR values associated with the $T$ received packets, we denote as $\hat{\gamma}_1, \hat{\gamma}_2, \ldots, \hat{\gamma}_T$ the ordered sequence of the same SNR values, such that $\hat{\gamma}_1 \geq \hat{\gamma}_2 \geq \cdots \geq \hat{\gamma}_T$. Taking into account that the probability density function (pdf) of each SNR value is given by $g(\hat{\gamma}_\lambda) = e^{-\gamma_\lambda/\overline{\gamma}}/\overline{\gamma}$, where $\lambda = 1, 2 \ldots, T$, we can obtain the joint pdf of the ordered SNR values using the properties of ordered statistics. In particular,

$$g(\hat{\gamma}_1, \ldots, \hat{\gamma}_T) = T! \prod_{\lambda=1}^{T} g(\hat{\gamma}_\lambda). \tag{16}$$

The probability of recovering the $\ell$ "strongest" packets, denoted as $q_\ell^T$, can be derived by integrating the joint pdf in (16), provided that the limits of integration for each individual SNR value are properly defined. Using the definition of the achievable rate $\mathcal{R}$ as the criterion for successful ($\mathcal{R} \geq 1$) or unsuccessful ($\mathcal{R} < 1$) packet recovery [5]

$$\mathcal{R} = C\left(\frac{\hat{\gamma}_\lambda}{\sum_{i=\lambda+1}^{T} \hat{\gamma}_i + 1}\right) \begin{cases} \geq 1, & \text{for } 1 \leq \lambda \leq \ell \\ < 1, & \text{for } \ell < \lambda \leq T, \end{cases} \tag{17}$$

where $C(z) \triangleq \log_2(1+z)$, we can express $q_\ell^T$ as follows

$$\begin{aligned} q_\ell^T = T! &\int_{\sum_{i=2}^{T} \hat{\gamma}_i + 1}^{+\infty} g(\hat{\gamma}_1) d\hat{\gamma}_1 \cdot \ldots \cdot \int_{\sum_{i=\ell+1}^{T} \hat{\gamma}_i + 1}^{+\infty} g(\hat{\gamma}_\ell) d\hat{\gamma}_\ell \\ &\cdot \int_{\hat{\gamma}_{\ell+2}}^{\sum_{i=\ell+2}^{T} \hat{\gamma}_i + 1} g(\hat{\gamma}_{\ell+1}) d\hat{\gamma}_{\ell+1} \cdot \int_{\hat{\gamma}_{\ell+3}}^{+\infty} g(\hat{\gamma}_{\ell+2}) d\hat{\gamma}_{\ell+2} \\ &\cdot \ldots \cdot \int_{\hat{\gamma}_T}^{+\infty} g(\hat{\gamma}_{T-1}) d\hat{\gamma}_{T-1} \cdot \int_0^{+\infty} g(\hat{\gamma}_T) d\hat{\gamma}_T. \end{aligned} \tag{18}$$

Computation of all integrals yields (14). ∎

In order to facilitate our analysis, we assume that the memory of all relays is full and, thus, the probability that a relay will randomly select a packet and forward it to the destination is $1/\nu$. Having derived an expression for $q_\ell^T$, we now proceed to the computation of $f(k_1, k_2)$.

*Proposition 2:* The probability that the destination will fail to recover a particular packet, when $k_1$ relays possess it but no more than $k_2 \leq k_1$ of those relays have transmitted it, is given by

$$\begin{aligned} f(k_1, k_2) = &\sum_{m=0}^{M-k_1} \binom{M-k_1}{m} p_L^{M-k_1-m} (1-p_L)^m \\ &\times \sum_{r=0}^{k_2} \binom{k_2}{r} \left(\frac{1}{\nu}\right)^r \left(1 - \frac{1}{\nu}\right)^{k_2 - r} \\ &\times \sum_{\ell=0}^{m+k_2} \frac{\binom{m+k_2-r}{\ell}}{\binom{m+k_2}{\ell}} q_\ell^{m+k_2}. \end{aligned} \tag{19}$$

*Proof:* The probability $f(k_1, k_2)$ can be computed by taking the average with respect to $m$, $r$ and $\ell$, where $m$ is the number of transmitting (thus, interfering) relays that do *not* possess the packet of interest in their memories, $r$ is the number of relays that transmit the packet of interest and $\ell$ is the number of successfully recovered packets. The fraction in the last row of (19) is the probability that the packet of interest is *not* among the $\ell$ recovered packets. ∎

## V. Results

We note that the model described in Section III does not consider the correlation between life cycles of neighboring packets. However, the ergodic principle grants that the average behavior of a packet can be well predicted by our model, provided that sufficiently long transmission bursts occur.

Let us consider a network with $M = 2$ relays, each having a memory size of $\nu = 1$. The source encodes $K = 100$ packets. The relationship between the erasure probability and the probability of listening, $p_L$, is shown in Fig. 3 for various relay-to-destination SNR values, $\overline{\gamma}$. Solid lines were obtained using our theoretical model; circles correspond to simulation measurement of the erasure probability over all packets in a realization, which is then averaged over all realizations; squares correspond to simulation measurement of the erasure probability for a given packet over all realizations, which is then averaged over all packets. Observe that simulation results are very close to theoretical predictions, confirming the ergodic nature of the system.

As another example, consider the case when $M = 3$, $\nu = 10$ and $K = 100$. As we can see in Fig. 4, the discrepancy between theoretical and simulation results is now more appreciable, although the matching is still acceptable. In this example, the value of $\nu$ is closer to $K$, which implies a stronger correlation between the life cycles of neighboring packets in a single realization. We can see in both Fig. 3 and Fig. 4 that there is an optimal value for $p_L$, denoted as $p_L^*$, which minimizes the erasure probability. The value of $p_L^*$ depends weakly on the relay-to-destination SNR and lies slightly below 0.5.

Having observed a good matching between simulations and theoretical results, we now use our theoretical model to explore the dependence of the minimum achievable erasure probability, which is observed when $p_L = p_L^*$, on the number of relays and the memory size. In Fig. 5, the minimum erasure probability is depicted as a function of the memory size, when $M = 3$ relays.

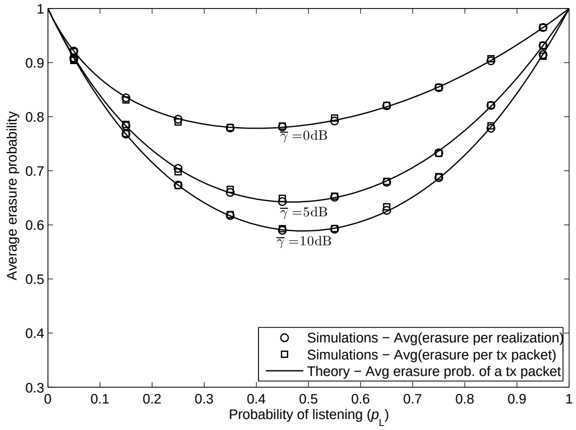


Fig. 3. Theoretical and simulation results for the erasure probability as a function of the probability of listening. Two relays and a unitary memory size are considered. For the simulations, $K=100$ source packets are encoded.

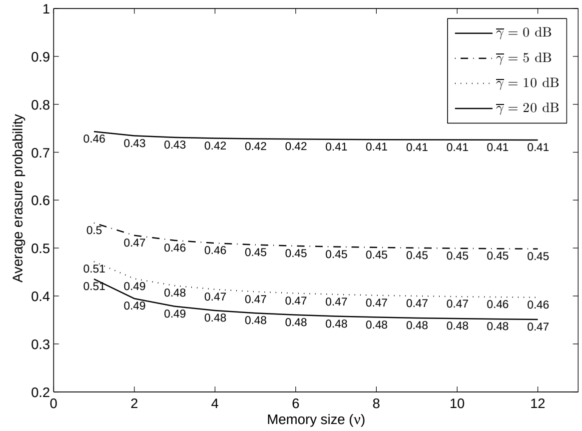


Fig. 5. Theoretical curves depicting the relationship between the minimum achievable erasure probability and the memory size at the relays. A network of three relays is considered. The values of $p_L^*$ appear on the curves.

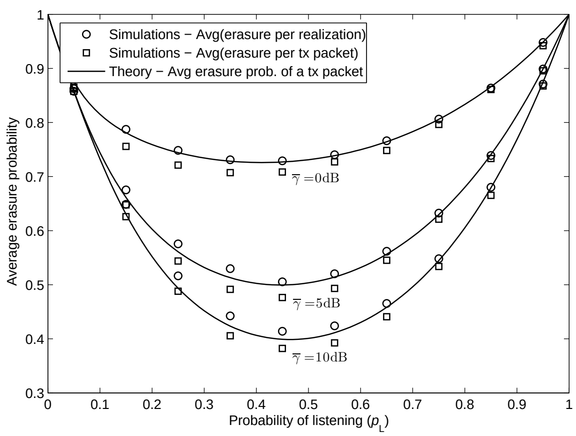


Fig. 4. Theoretical and simulation results for the erasure probability as a function of the probability of listening. Three relays and a memory size of ten are considered. For the simulations, $K=100$ source packets are encoded.

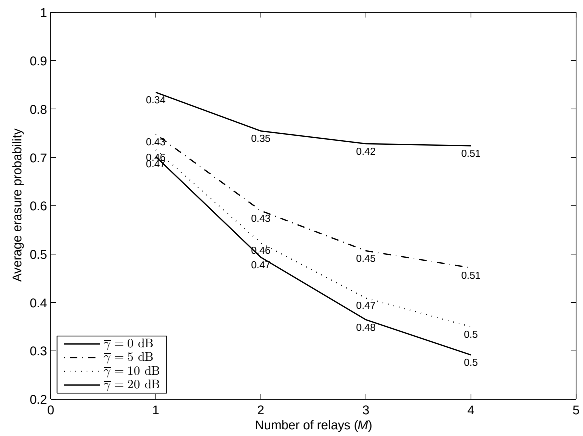


Fig. 6. Theoretical curves depicting the relationship between the minimum achievable erasure probability and the number of relays in the network. A memory size of five is considered. The values of $p_L^*$ appear on the curves.

Various values for $\overline{\gamma}$ are considered, while the value of $p_L^*$ is also reported on the curves. It can be seen that the erasure probability slightly decreases with $\nu$, but reaches a constant value as $\nu$ increases. We observed a similar behavior for different values of $M$. The relationship between the minimum erasure probability and the number of relays, for a memory size of $\nu=5$, is presented in Fig. 6. Comparing the two figures, we note that changes to the number of relays have a greater impact on the erasure probability than changes to the memory size. Note, however, that the slope of the curves in Fig. 6 gradually diminishes as $M$ increases.

## VI. Conclusion

In this paper, we presented a model for single-hop relay networks, in which relays randomly select coded packets that were previously received, CRC validated and stored in their memory, and forward them to the destination. Using a semi-analytical method based on Markov chains, we demonstrated that the erasure probability of the system can be minimized by selecting proper values for the probability of listening, the number of relays and the memory size at each relay.

Note that the proposed scheme can be seen as an instance of distributed fountain codes; its generalization, in which relays can also XOR packets, will be carried out in future work.

## Acknowledgment

This work was supported by EPSRC, Politecnico di Torino, the University of Cambridge, NEWCOM++ and the Royal Academy of Engineering.

## References

[1] J. Byers, M. Luby, M. Mitzenmacher, and A. Rege, "A digital fountain approach to reliable distribution of bulk data," in *Proc. ACM SIGCOMM*, Vancouver, Canada, Sept. 1998, pp. 56–67.

[2] D. J. C. MacKay, "Fountain codes," *IEE Proceedings*, vol. 152, no. 6, pp. 1062–1068, Dec. 2005.

[3] A. F. Molisch, N. B. Mehta, J. S. Yedidia, and J. Zhang, "Performance of fountain codes in collaborative relay networks," *IEEE Trans. Wireless Commun.*, vol. 6, no. 11, pp. 4108–4119, Nov. 2007.

[4] J. Castura and Y. Mao, "Rateless coding for wireless relay channels," *IEEE Trans. Wireless Commun.*, vol. 6, no. 5, pp. 1638–1642, May 2007.

[5] H. Wicaksana, S. H. Ting, and Y. L. Guan, "Spectral efficient half duplex relaying for fountain code with wireless network coding," in *Proc. IEEE Int. Conf. Commun.*, Beijing, China, May 2008, pp. 295–299.

[6] C. M. Grinstead and J. L. Snell, *Introduction to Probability*, 2nd ed. American Mathematical Society, 1997, ch. 11.