

Delay Constrained Throughput-Reliability Tradeoff in Network-Coded Wireless Systems

David C. Adams, Jinfeng Du*, and Muriel Médard
Massachusetts Institute of Technology
Cambridge, 02139, MA
Email: {dcadams, jinfeng, medard}@mit.edu

Christopher C. Yu
Draper Laboratory
Cambridge, 02139, MA
Email: cyu@draper.com

Abstract—We investigate the performance of delay constrained data transmission over wireless networks without end-to-end feedback. Forward error-correction coding (FEC) is performed at the bit level to combat channel distortions and random linear network coding (RLNC) is performed at the packet level to recover from packet erasures. We focus on the scenario where RLNC re-encoding is performed at intermediate nodes and we assume that any packet that contains bit errors after FEC decoding can be detected and erased. To facilitate explicit characterization of data transmission over network-coded wireless systems, we propose a generic two-layer abstraction of a network that models both bit/symbol-level operations at the lower layer (termed PHY-layer) over several heterogeneous links and packet-level operations at the upper layer (termed NET-layer). Based on this model, we propose a network reduction method to characterize the throughput-reliability function of the end-to-end transmission. Our approach not only reveals an explicit tradeoff between data delivery rate and reliability, but also provides an intuitive visualization of the bottlenecks within the underlying network. We illustrate our approach via a point-to-point link and a relay network and highlight the advantages of this method over capacity-based approaches.

Index Terms—Wireless networks, random linear network coding, delay, throughput, reliability, cross-layer optimization

I. INTRODUCTION

In wireless communication systems, much effort has been devoted to advanced forward error-correction coding (FEC) and signal processing techniques on the physical layer to pursue reliable transmission over each single hop. Even if data transmission over each individual link is reliable, the end-to-end transmission can still fail as a result of packet losses caused by higher-layer effects, such as congestion or buffer overflow. These missing packets are usually detected and recovered by some automatic repeat-request (ARQ) based mechanisms on data link layer and/or on transport layer. Many wireless networks are inherently heterogeneous in the sense that the channel quality, system implementation, and available resources can be vastly different from link to link. These parameters may interact with and be dependent upon one another, which can make the performance analysis a difficult task over networks of nontrivial size.

In this paper, we focus on delay constrained data transmission over wireless networks where end-to-end feedback is absent or performed at a higher layer, motivated by scenarios where the end-to-end feedback is unavailable, excessively delayed, or onerous. While, from solely a throughput perspective,

*Jinfeng Du is also with Royal Institute of Technology, Stockholm, Sweden.

coding at the physical layer should suffice if the only losses are due to deleterious effects at that layer, operational wireless communication systems differ from this paradigm in operation. Firstly, they are comprised of at least two layers, which we coarsely identify as PHY and NET; the latter operates on the packet level. Secondly, some losses are not due to physical layer effects (such as buffer overflow). Finally, throughput may not be the only parameter we seek to improve; in particular, we consider the tradeoff between throughput and reliability, under a delay constraint.

To model this behavior, we propose a generic two-layer abstraction of the wireless network incorporating the effects of bit/symbol-level operations at the PHY over multiple heterogeneous links and packet-level operations at the NET. Our model assumes that packets are either provided to the NET layer intact or erased, a simplification based on the operations practiced by state-of-the-art systems. PHY level techniques (such as adaptive coding and modulation), while imperative in achieving this effect, are not explicitly considered here. We focus on the scenario where intermediate nodes can perform random linear network coding (RLNC) [1] re-encoding based on the received packets. The interaction of parameters from the different layers is characterized by a clean and general interface; we assume that we always know the PHY-layer operations. Based on this model, we propose a network reduction method that works on the NET-layer of the wireless system to characterize the throughput-reliability function of the end-to-end transmission under the delay constraint. Our approach not only reveals an explicit tradeoff between data delivery rate and reliability, but also provides an intuitive visualization of the bottlenecks within the underlying network. We illustrate our approach via a point-to-point link and a relay network and highlight its advantages over capacity-based approaches.

The rest of this paper is organized as follows. We first give a brief overview of related work in Sec. II and then present the two-layer model in Sec. III. We analyze the probability of errors on both layers in Sec. IV and describe the network reduction method in Sec. V. In Sec. VI we establish the throughput-reliability function and illustrate via examples how it can be used for system design. Sec. VII is the conclusion.

List of Notation

- \mathbb{F}_q^l : set of all length- l vectors over the finite field \mathbb{F}_q

- $\bar{\lambda} \triangleq [\bar{\lambda}[0], \dots, \bar{\lambda}[n]] \in [0, 1]^{n+1}$: a vector representing the probability mass function evaluated at $\{0, 1, \dots, n\}$
- $\bar{\Lambda}^c$: a vector representing the complementary cumulative distribution function (i.e., $\bar{\Lambda}^c[k] = \sum_{i=k}^n \bar{\lambda}[i]$)
- $*$: operator for convolution of two vectors
- \odot : element-wise multiplication (Hadamard product).

II. RELATED WORK

Numerous research efforts have been devoted to characterizing the interaction between channel coding and packet-erasure coding from different aspects; we only list a sampling here. Vehkaperä and Médard in [2] propose a two-layer concatenated code framework with a total delay constraint, where error exponents for random coding are used to analyze both inner and outer codes. This framework investigates a throughput-delay tradeoff, clarifies the embedding of error functions in a multi-level approach, and models the fact that data may be lost at either layer. Barg et al. in [3] also consider error exponents, although the focus is on joint channel and packet coding rather than decoupled codes and the regime is asymptotic in block lengths at both layers. Berger et al. in [4] focus on a flat block-fading channel, where perfect channel decoding is assumed for all rates below capacity and perfect packet-erasure coding is assumed when a fixed overhead requirement is met. The tradeoff between spectral efficiency and end-to-end probability of error is explored; one particularly useful insight is the investigation of this tradeoff with few choices of physical layer transmission modes. Courtade and Wesel in [5] also assume perfect channel and erasure coding as in [4], but consider both fast and slow flat block-fading with a delay constraint. They employ a Gaussian approximation of end-to-end outage probability and optimize for power consumption under end-to-end rate and reliability constraints. In [6], Koller et al. consider RLNC operating on top of a random code over a binary symmetric channel (BSC) for a one-hop broadcast network. The expected number of transmissions required to decode is minimized, and it is noted that this optimization problem is not equivalent to optimizing the expected number of accurate bits per transmission.

While [2]–[6] all elucidate different angles of the cross-layer problem, they all consider only single-link networks where generalization to larger topologies is nontrivial. We also recognize some works that consider more complicated network structures. Swapna et al. in [7] work on RLNC over k packets broadcasting to n users over independent time-correlated erasure channels. They conclude that k must scale at $O(\log(n))$ to ensure non-zero throughput. Under such conditions, the mean and variance of total transmission time depend on channel correlation. Lun et al. in [8] propose a framework to translate a lossy unicast/multicast network into a lossless packet network with only point-to-point arcs (for wireline) and hyper-arcs (for wireless). Assuming that the number of packets is large and that the arrivals of packets at each node are independent Poisson processes, the probability of RLNC decoding error is characterized by the delay, rate, and the network capacity. The average throughput of each individual link is the figure of merit

characterizing system performance [8]; additionally, delay is based entirely on the depth of the network independently of propagation time. Lower layers are not considered in [8]. Ming et al. in [9] investigate the delay in packet erasure networks where RLNC is used in a rateless fashion, and the delay is optimized based on the tradeoff between codeword lengths on physical layer and on network layer.

III. SYSTEM MODEL: TWO-LAYER ABSTRACTION

A source node intends to transmit some messages to a remote destination via a wireless network. The transmission has to be completed within a certain delay constraint with a rate of success no less than a predefined threshold. The underlying wireless network connecting the source and the destination suffers from both deleterious effects on the PHY-layer and packet losses at the NET. The two layer abstraction of the source node is shown in Fig. 1.

A. Abstraction of the NET-Layer

On the NET layer, a new generation of source messages is first evenly split into κ segments $V_1, V_2, \dots, V_\kappa$, each of ηl bits (appending zeros when necessary) where $0 < \eta < 1$. These message segments are fed to the RLNC encoder to generate n RLNC-coded messages of the same size, which are then passed through the packetization process where a header of $(1-\eta)l$ bits is added to each of the RLNC-coded messages to formulate packets of size l bits. The above process can be summarized as follows

$$(\mathbb{F}_2^{\eta l})^\kappa \xrightarrow{\text{RLNC}} (\mathbb{F}_2^{\eta l})^n \xrightarrow{\text{packetization}} (\mathbb{F}_2^l)^n. \quad (1)$$

Here we assume that the packet length l and format (overhead) are long-term, system-wide parameters that cannot be adjusted based on a single transmission task. These RLNC-coded packets are then passed to the PHY layer for transmission. After a period of time τ , if the destination receives at least κ packets and the decoding is successful, we say the transmission task is accomplished. Otherwise the task fails. The benefits of choosing RLNC as the erasure code are two fold: its decoding error probability is easy to analyze, and its re-encoding process at intermediate nodes is straightforward and facilitates the network reduction approach we propose in this paper.

Our NET-layer model is generic in the sense that we do not refer to any specific network protocols. Rather, we use the abstraction of a “packet” as the atomic unit of information that is exchanged between the two layers. Note that in canonical Internet layering models, a packet has different meaning in different layers: a packet in an upper layer becomes the payload of another packet in the adjacent lower layer. In our model, a packet $[X_k, O_k]$ differs from the canonical concept, since its payload X_k refers to the RLNC-coded message (cf. service data unit on transport layer), but its overhead O_k refers to all bits other than the payload: all the control/protocol information (cf. overhead brought by data link layer, network layer, and transport layer) plus the RLNC encoding coefficients. The proportion of the payload (η) within each packet is a system parameter that reflects the actual network protocol in use.

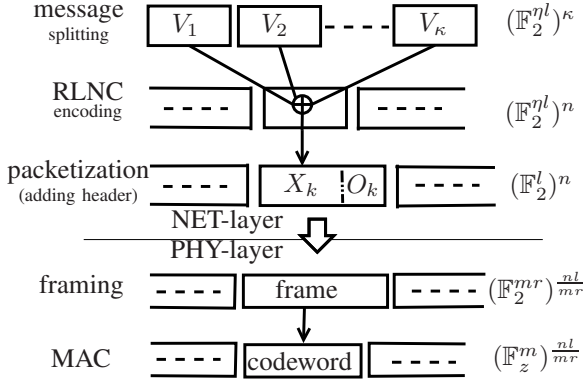


Fig. 1. Two layer abstraction of the source node. At the NET, a new generation of source message is first evenly split into κ segments $V_1, V_2, \dots, V_\kappa$, each of ηl bits (appending zeros when necessary), and then passed through RLNC encoding (generating payload X_k) and packetization process (adding header O_k) to generate packets of size l bits. These packets are passed to the PHY layer where they are grouped into frames of mr bits, and then translated into m -symbol codewords over \mathbb{F}_z^m by operations such as the channel coding and modulation for media access (abbr. MAC).

B. Abstraction of the PHY-layer

On the PHY-layer, packets from the NET-layer are grouped into frames¹ of mr bits, and each frame is then mapped onto an m -symbol sequence (a.k.a., codeword) over \mathbb{F}_z^m by operations such as the channel coding and modulation process for media access (abbr. MAC). The operation at the PHY can be described by the following mapping

$$(\mathbb{F}_2^l)^n \rightarrow (\mathbb{F}_2^{mr})^{\frac{nl}{mr}} \rightarrow (\mathbb{F}_z^m)^{\frac{nl}{mr}}. \quad (2)$$

Since each codeword carries a message of mr bits, or equivalently, $\alpha = mr/l$ packets, we call α the packet loading parameter as it decouples the PHY and the NET cleanly such that a packet is the basic unit at the NET and a message/codeword is the basic unit at the PHY.

The MAC block at the PHY does not specify or rely on any specific coding and modulation techniques. In effect, any physical layer operations that affect the transition from frames to codewords can be incorporated into the MAC block by adjusting the data rate r (bits per symbol-time, or equivalently, bits per channel-use). By focusing on r at the PHY, our model facilitates flexible MAC, channel coding, and modulation schemes in the system. For example, when two users share a channel via hierarchical modulation where z -QAM is used, and their channel coding rate are c_1 and c_2 , respectively, the effective data rates are $r_i = \frac{1}{2}c_i \log_2(z)$, $i = 1, 2$. The corresponding MAC operation at the PHY is described by

$$(\mathbb{F}_2^{mr_1}, \mathbb{F}_2^{mr_2}) \rightarrow (\mathbb{F}_2^{\frac{r_1}{c_1}m}, \mathbb{F}_2^{\frac{r_2}{c_2}m}) \rightarrow (\mathbb{F}_{\sqrt{z}}^m, \mathbb{F}_{\sqrt{z}}^m) \rightarrow \mathbb{F}_z^m. \quad (3)$$

C. Interface between Two Layers

Based on the two-layer abstraction, we can model and track the data flow over wireless networks as illustrated in Fig. 2 for transmission over a two-hop network. Note that at the

¹Note that our framing operation differs from canonical packet aggregation/segmentation in the sense that it incurs no change on the structure of a packet, which ensures independent packet erasures at a cost of overhead.

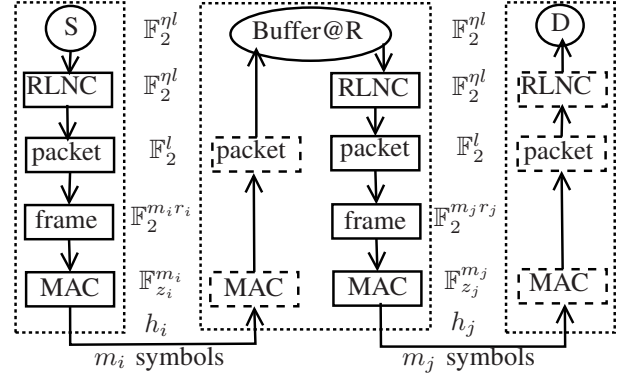


Fig. 2. Flow chart of a two-hop network where the source sends m_i -symbol codewords at data rate r_i over channel h_i to the relay. The relay stores its received packets in a buffer, performs RLNC over the buffer, and then forwards the re-encoded packets to the destination over channel h_j with PHY layer parameters (m_j, r_j) . Dashed blocks indicate the decoding process. The elements in each stage are represented by vectors over finite fields.

relay a new RLNC-coded packet is generated in two steps: its payload is a random linear combination of the payloads of all the received RLNC packets stored in its buffer², and the coding vector in the header is likewise the corresponding linear combination of coding vectors in the buffer.

At the PHY, there are many dependent operating parameters even on a single link i : channel h_i , codeword length m_i , data rate r_i , channel coding rate c_i , modulation order z_i , packet loading factor α_i , and codeword error probability $\tilde{p}_{e,i}$. Regardless of the PHY-layer operations, we only focus on two parameters that directly affect the NET: n_i (number of packets transmitted within the delay constraint τ) and $p_{e,i}$ (packet erasure probability, which may be different from the codeword error probability $\tilde{p}_{e,i}$ as discussed in Sec. IV).

At the NET, the number of RLNC coded packets that can be transmitted over a channel is limited by the PHY-layer settings. A packet can be either erased at the PHY with probability $p_{e,i}$, or dropped by the NET-layer effects with probability $p_{t,i}$. We denote the overall packet erasure probability as a function

$$\xi_i = Pe(p_{e,i}, p_{t,i}, l). \quad (4)$$

Now for each link, at the NET we only focus on n_i and ξ_i (the overall packet erasure probability).

IV. ERROR ANALYSIS

Given a generation of κ messages $W = \{V_1, V_2, \dots, V_\kappa\}$ each of ηl bits, by representing them as vectors over \mathbb{F}_q , a RLNC-coded packet X_k is constructed by

$$X_k = a_{k,1}V_1 + a_{k,2}V_2 + \dots + a_{k,\kappa}V_\kappa, \quad (5)$$

where the encoding coefficients $\{a_{k,1}, a_{k,2}, \dots, a_{k,\kappa}\}$ are drawn uniformly at random from \mathbb{F}_q . The re-encoding process at intermediate nodes is the same as in (5) except that the linear combination is over the payloads of received packets. By the end of the delay constraint τ , some of the coded packets will have been received together with their encoding vectors.

²As pointed out in [10], the buffer can be as small as one packet in size.

If these encoding vectors formulate a matrix of rank κ , we are able to recover W through Gaussian elimination. If the rank is less than κ , the transmission fails. Through any link, a packet may be erased with probability ξ , which depends both on local parameters p_e and p_t as well as the packet length l (a system parameter).

At the PHY, a codeword consisting of m symbols is transmitted across the channel. The corresponding decoding error probability \tilde{p}_e depends on many parameters as explained in Sec. III-C. We may approximate the error probability by error exponent bounds as in [2], [3], by large deviation methods, or by the recent finite block length bounds developed in [12]. Alternatively, we can always emulate the channel and draw the rate-error probability curve based on engineering practice. Due to the framing process at the PHY, one codeword contains $\alpha=mr/l$ packets. If $\alpha=1$, i.e., one packet per codeword, the packet erasure probability p_e at the PHY equals the codeword error probability \tilde{p}_e . If one codeword contains several packets, each packet can still experience independent losses owing to their independent cyclic redundancy check (CRC) embedded in the packet headers. For the case $\alpha < 1$ or non-integer α , we have $p_e > \tilde{p}_e$ and their relationship depends on the actual framing process. It is therefore favorable to have integer-valued α in system design.

Let S be the number of packets successfully received by the end of transmission, and let D be the number of degrees of freedom (DoF) available at the decoder (i.e., the rank of the matrix composed by encoding vectors). Given i.i.d. packet erasure with probability $\xi \in [0, 1]$, the probability of success of the RLNC coded transmission can be written as (see [8], [11])

$$Pr\{D=\kappa\} = \sum_{s=\kappa}^n \binom{n}{s} (1-\xi)^s \xi^{n-s} \prod_{u=0}^{\kappa-1} (1-q^{u-s}). \quad (6)$$

This formulation is exact, but it is also difficult to precisely compute when $\frac{\kappa}{n}$ is well below the channel capacity $(1-\xi)$. With the help of an inequality given by Liva et al. in [11],

$$1 - \frac{1}{q-1} q^{\kappa-s} < \prod_{u=0}^{\kappa-1} (1-q^{u-s}) \leq 1 - q^{\kappa-s-1}, \quad (7)$$

we can get a lower bound for probability of error as follows

$$\begin{aligned} P_e(\kappa) &= 1 - Pr\{D = \kappa\} \\ &\geq 1 - \sum_{s=\kappa}^n \binom{n}{s} (1-\xi)^s \xi^{n-s} (1 - q^{\kappa-s-1}). \end{aligned} \quad (8)$$

This lower bound is very tight when q is large.

V. NETWORK REDUCTION

In the following discussion we assume that our RLNC re-encoding operations are non-degenerate (i.e., q is sufficiently large); if the coding vectors of the packets received at any node span a subspace \mathcal{V} of \mathbb{F}_q^κ , where κ is number of information messages within a generation over which RLNC is performed, then any κ of the new vectors of the re-encoded packets also span \mathcal{V} . In such cases, the probability of success defined in (6) is essentially the complementary cumulative distribution function (CCDF) evaluated at $S = \kappa$.




Fig. 3. Basic two-link parallel network and two-link tandem network.

A. Tandem Link Equivalence

Let our network consist of two directed links connecting three nodes, as in the right part of Fig. 3. We will call these links 1 and 2. Link $i \in \{1, 2\}$ permits us to send n_i packets per delay period τ . Now define a random variable S_i describing the number of successfully delivered packets over the link i within the generation. Assuming that we have an explicitly defined channel model, we can get a probability mass function (PMF) for S_i , which will be described by the vector $\bar{\lambda}_i \in [0, 1]^{n_i+1}$, where

$$\bar{\lambda}_i[s] = Pr\{S_i = s\}, \quad s \in \{0, 1, \dots, n_i\}, \quad (9)$$

or equivalently by the CCDF of S_i , $\bar{\Lambda}_i^c \in [0, 1]^{n_i+1}$:

$$\bar{\Lambda}_i^c[s] = Pr\{S_i \geq s\} = \sum_{j=s}^{n_i} \bar{\lambda}_i[j], \quad s \in \{0, 1, \dots, n_i\}. \quad (10)$$

Because our assumption permits us to ignore the possibility of a degenerate decoding matrix, the number of DoF D at the tail of the second link is equal to the minimum number of packets received across either link. For the distribution of D , we note that the following property holds for $\kappa \leq \min\{n_1, n_2\}$,

$$\begin{aligned} Pr\{D \geq \kappa\} &= Pr\{\{S_1 \geq \kappa\} \cap \{S_2 \geq \kappa\}\} \\ &= Pr\{S_1 \geq \kappa\} Pr\{S_2 \geq \kappa\}, \end{aligned} \quad (11)$$

where the second equality follows from the independence of the two links. We have now found the CCDF of a single link equivalent to that of the two-link tandem network. Then, if we zero-pad the shorter of the two single-link CCDFs and denote this operation with a $'$, we have

$$\bar{\Lambda}_{EQ}^c[d] = \bar{\Lambda}'_1[d] \bar{\Lambda}'_2[d], \quad (12)$$

or equivalently,

$$\bar{\Lambda}_{EQ} = \bar{\Lambda}'_1 \odot \bar{\Lambda}'_2, \quad (13)$$

where \odot denotes element-wise multiplication of two vectors.

B. Parallel Link Equivalence

Now we consider the scenario of two links in parallel, as shown in the left part of Fig. 3. This problem can be regarded as a kind of a dual to the tandem link problem. Let n_i and S_i retain their definitions from earlier. If we expand our assumption so that packets received from either link are linearly independent, the DoF D at node B becomes

$$D = S_1 + S_2. \quad (14)$$

Because S_1 and S_2 are independent, we can get a distribution for D by convolving the distributions of S_1 and S_2 :

$$\bar{\lambda}_{EQ} = \bar{\lambda}_1 * \bar{\lambda}_2, \quad \bar{\lambda}_{EQ} \in [0, 1]^{n_1+n_2+1}. \quad (15)$$

Notice that we can quite easily switch between $\bar{\lambda}$ and $\bar{\Lambda}^c$ using a 1-to-1 linear operator.




Fig. 4. The throughput ($\kappa \times l / \tau$ bits per second)-reliability ($P_e(\kappa)$) functions of two erasure channels with the same capacity, given packet length l and delay constraint τ . CCDF and PMF are also plotted for references.

VI. THROUGHPUT-RELIABILITY TRADE-OFF

We inherit the setting of Lun et al. in [8], where a probability of error is obtained based on long-term average. By reducing a simple network to one distribution for received packets, given by $\bar{\lambda}$ or $\bar{\Lambda}^c$, we characterize a throughput-reliability tradeoff for each generation of data under time constraint τ . The expected number of DoF can be calculated by summing $\bar{\Lambda}^c \in [0, 1]^{n+1}$,

$$E[D] = \sum_{d=1}^n d\bar{\lambda}[d] = \sum_{d=1}^n \bar{\Lambda}^c[d] = \sum_{d=0}^n \bar{\Lambda}^c[d] - 1 \quad (16)$$

We can prove that $E[D]$ asymptotically approaches the cut-set capacity as $\tau \rightarrow \infty^3$. Furthermore, $1 - \bar{\Lambda}^c[\kappa]$ produces the probability of error for any choice of κ , making it an attractive option for system design. Since the effective rate is $\frac{\kappa l}{\tau}$ bits per second, $1 - \bar{\Lambda}^c[\kappa]$ can be regarded as the throughput-reliability function given any specific l .

The advantage of our approach over capacity based approaches (e.g., as in [8]) can be visualized in Fig. 4 where two links have the same capacity but differ largely in throughput-reliability functions. With the throughput-reliability curve, system designers have the option of selecting objective functions more informative than throughput alone. The throughput-reliability curve also facilitates the identification of the maximum generation size for a minimum reliability constraint or the minimum probability of error for a given generation size. Also, the end-to-end CCDF makes it easy to identify the bottleneck in a network; by condensing large pieces of a

³Due to space limitation, the proof can be found in [13].




Fig. 5. Results of optimization for $\kappa(1 - P_e(\kappa))$ on a single link with one modulation scheme and choice of 7 RS codes producing a range of n from 20 to 140. The underlying channel BSC(ϵ) varies on the abscissa, while p_t is fixed at 0.1. Optimal RS coding rate is shown in (a), NET-layer coding rate (κ/n) in (b), and the goodput $\kappa(1 - P_e(\kappa))$ in (c).

network to a few resultant distributions, we can pick out the challenged area as the CCDF that tapers off the fastest.

A. Illustrations

For the purpose of illustration, we assume all PHY-layer symbols consume the same transmission time τ_s . Given the delay constraint τ , the maximum number of symbols that can be transmitted on the NET layer is bounded by

$$n_s \leq (\tau - 0.5\tau_{RTT})/\tau_s,$$

where τ_{RTT} is the round trip time.

In Fig. 5, we show the theoretical results of code optimization over a practically-motivated scenario for one link where the channel at the PHY is a BSC(ϵ) where ϵ is the bit-cross-over probability. We envision a Reed-Solomon (RS) code with fixed block length $m=255$ coding symbols (each of one-byte) and only 7 choices for coding rate. Over this, we run RLNC over \mathbb{F}_{2^8} (i.e., one-byte encoding coefficients) where the number of RLNC-coded packets n is limited by the delay constraint and the RS coding rate. While varying bit-cross-over probability on the underlying channel and fixing p_t , we evaluate $\kappa(1 - P_e(\kappa))$, the “goodput”, for every configuration of (r, κ) , recording the highest value. The top line shows the optimal RS code rate, the second line shows the optimal RLNC rate, and the bottom line shows the goodput. Interestingly, we see that the network code rate drops down when the RS code begins to fail and erasures due to decoding on the lower level become more prevalent than dropped packets. This creates a smoothing effect on throughput.

A similar setup is used on the relay network in Fig. 6 to generate results for Fig. 7; each link has the same selection of RS codes with slightly different, but fixed, channel parameter and NET packet erasure rate. To simplify simulation we assume that $\xi_i = p_{e,i} + p_{t,i}$. With the chosen BSC channel parameters $(\epsilon_1, \epsilon_2, \epsilon_3)$ as shown in Fig. 6, the corresponding RS codeword error probability $p_{e,1}=0.1$ at RS rate $7/8$, $p_{e,2}=0.25$ at rate




Fig. 6. A relay network where link $i \in \{1, 2, 3\}$ has a BSC(ϵ_i) at the PHY and a packet erasure probability p_{ti} at the NET.

4/8, and $p_{e,3} = 0.32$ at rate 3/8. The figure shows four different optimization results. The stepped curve results (solid line) is obtained by first selecting an end-to-end reliability constraint and then finding the configuration whose end-to-end CCDF (obtained using the network reduction technique described in Sec. V) provides the highest rate under this reliability constraint. The curve (with marker \times) for the expected end-to-end DoF (i.e., the capacity) is obtained based on the end-to-end CCDF after network reduction. We also optimize the expected DoF to cross each link individually, and then calculate the end-to-end CCDF based on the local optimal configurations. The result is indicated by the curve with marker \circ . Finally, we have one point (\blacksquare) indicating the maximum possible value of the goodput $\kappa(1 - P_e(\kappa))$. Note that optimizing for expected number of DoF does not necessarily yield the same configuration as optimizing the rate given a reliability constraint.

VII. CONCLUSION

In this paper we have explored the relationship between FEC on the PHY layer and RLNC on the NET layer over simple network flows with an end-to-end delay constraint. Working from first principles, we developed an abstract, skeletal model for generality, then applied the properties of RLNC to obtain more detailed results. We briefly introduced a way to combine links on the NET layer so as to derive the distribution of the amount of information received per generation, instead of only the expected amount (capacity). We argued that this distribution can guide us to more efficient transmissions than capacity alone, and made an effort to convey this point through some practically-motivated optimization problems.

This work presents opportunity for extension on a few fronts. First, we have developed in [13] an argument for the validity of our network reduction procedure as well as a convex optimization approach for system design. Furthermore, our reduction operations cannot reduce arbitrary network flows in their current state and new approaches have been developed in [14] to provide very tight upper and lower bounds for the end-to-end throughput-reliability curves. Finally, our abstraction could be useful in contexts other than constrained delay.

ACKNOWLEDGMENT

This work is funded in part by the Swedish Research Council (VR) and by VT iDirect. This material is based upon work supported by the Air Force Office of Scientific Research (AFOSR) under award No. FA9550-13-1-0023.




Fig. 7. Performance of the relay network with limited PHY modes after different optimization procedures. The stepped curve indicates the maximum generation size (κ) while maintaining the error constraint on the abscissa. Also shown are (P_e, κ) pairs when the PHY rates are selected to maximize expected DoF across both the entire network and each link individually. Finally, the red square indicates the maximum goodput ($\kappa(1 - P_e(\kappa))$) achievable for any PHY configuration and choice of κ .

REFERENCES

- [1] T. Ho, M. Médard, R. Koetter, M. Effros, D. R. Karger, J. Shi, and B. Leong, "A random linear network coding approach to multicast," *IEEE Trans. Inf. Theory*, vol. 52, pp. 4413–4430, Oct. 2006.
- [2] M. Vehkaperä and M. Médard, "A throughput-delay trade-off in packetized systems with erasures," in *Proceedings IEEE ISIT*, Sep. 2005.
- [3] A. Barg, J. Justesen, and C. Thommessen, "Concatenated codes with fixed inner code and random outer code," *IEEE Trans. Inf. Theory*, vol. 47, pp. 361–365, Jan. 2001.
- [4] C. R. Berger, S. Zhou, Y. Wen, P. Willett, and K. Pattipati, "Optimizing joint erasure- and error-correction coding for wireless packet transmissions," *IEEE Trans. Wireless Commun.*, vol. 7, pp. 4586–4595, Nov. 2008.
- [5] T. A. Courtade and R. D. Wesel, "Optimal allocation of redundancy between packet-level erasure coding and physical-layer channel coding in fading channels," *IEEE Trans. Commun.*, vol. 59, pp. 2101–2109, Aug. 2011.
- [6] C. Koller, M. Haenggi, J. Kliewer, and D. J. Costello, "On the optimal block length for joint channel and network coding," in *Proceedings IEEE ITW*, Oct. 2011.
- [7] B. T. Swapna, A. Eryilmaz, and N. B. Shroff, "Throughput-delay analysis of random linear network coding for wireless broadcasting," *IEEE Trans. Inf. Theory*, vol. 59, pp. 6328–6341, Oct. 2013.
- [8] D. S. Lun, M. Médard, R. Koetter, and M. Effros, "On coding for reliable communication over packet networks," *Physical Communication*, vol. 1, pp. 3–20, 2008.
- [9] M. Xiao, M. Médard, and T. Aulin, "Cross-layer design of rateless random network codes for delay optimization," *IEEE Trans. Commun.*, vol. 59, pp. 3311–3322, Dec. 2011.
- [10] B. Haeupler and M. Médard, "One packet suffices - Highly efficient packetized network coding with finite memory," in *Proceedings IEEE ISIT*, Aug. 2011.
- [11] G. Liva, E. Paolini, and M. Chiani, "Performance versus overhead for fountain codes over \mathbb{F}_q ," *IEEE Commun. Letters*, vol. 14, pp. 178–180, Feb. 2010.
- [12] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Channel coding rate in the finite blocklength regime," *IEEE Trans. Inf. Theory*, vol. 56, pp. 2307–2359, Oct. 2013.
- [13] D. C. Adams, *A Delay-Constrained Cross-Layer Model Using Network Coding*, Master thesis, MIT, May 2014.
- [14] N. Sweeting, *Reduction of Arbitrary Networks: A Heuristic Approach*, RSI Report, MIT, Jul. 2014.