Throughput-Delay Analysis of Random Linear Network Coding for Wireless Broadcasting

Swapna B.T., Atilla Eryilmaz, and Ness B. Shroff
Departments of ECE and CSE
The Ohio State University
Columbus, OH 43210
Email: {buccapat, eryilmaz, shroff}@ece.osu.edu

***Abstract*—In this paper, we investigate the throughput and decoding-delay performance of random linear network coding as a function of the coding window size and the network size in an unreliable single-hop broadcast network setting. Our model consists of a source transmitting packets of a single flow to a set of N receivers over independent erasure channels. The source performs random linear network coding (RLNC) over K (coding window size) packets and broadcasts them to the receivers. We find that the broadcast throughput of RLNC must vanish with increasing N, for any given K. Hence, in contrast to other works in the literature, we investigate how the coding window size K must scale for increasing N. By appealing to the Central Limit Theorem, we approximate the Negative Binomial random variable arising in our analysis by a Gaussian random variable. We then obtain tight upper and lower bounds on the mean decoding delay and throughput in terms of K and N. Our analysis reveals that the coding window size of $\ln(N)$ represents a phase transition rate below which the throughput converges to zero, and above which it converges to the broadcast capacity. Our numerical investigations show that the bounds obtained using the Gaussian approximation also apply to the real system performance, thus illustrating the accuracy of the analysis.**

***Index Terms*—Broadcast, Delay Analysis, Erasure Channel, Network Coding.**

I. Introduction

We consider an important transmission scenario, occurring in many communication systems, whereby a source must broadcast common information to many receivers over wireless channels in a timely manner. Such a scenario occurs, for example, in a satellite or cellular network where a satellite or base station broadcasts a large file or streaming multi-media data to many receivers within their footprint over unreliable channels. Another example occurs in a multi-hop wireless network where each node broadcasts control information to all its immediate neighbors to coordinate medium access, power control, and routing operations. We note that such local sharing of control information (such as queue-length or other pricing information) is common to many provably efficient network controllers (e.g. [1], [2], [3] etc.).

In this work, the essential components of such wireless broadcast systems are modeled through a transmitter broadcasting consecutive blocks of K data packets over independently fading erasure channels with erasure probability p to N receivers. Assuming that the transmitter is infinitely backlogged, we consider transmission strategies that transfer the data in blocks of K packets, which include the class of block coding strategies. Among all such block transmission strategies, it has previously been shown (see [6]) that, for any fixed N and K, Random Linear Network Coding (RLNC) strategy (see Section III for a detailed description) asymptotically[1] minimizes the number of transmissions required to complete the transfer of all K packets at all N receivers (also called the *block completion time*).

With this motivation, we focus on the scaling performance of RLNC as a function of K and N with respect to the following two key metrics: the *(broadcast) throughput*, defined as the data transfer rate to all receivers; and the *(broadcast) decoding delay*, defined as the amount of time spent between the start of a block transmission and its completion (i.e. successful decoding) at all the receivers.

It is not difficult to see that the *(broadcast) capacity* of such a collection of N erasure channels, for any N, is equal to $(1 - p)$ packets per time slot. Moreover, this maximum limit on the throughput can be arbitrarily closely achieved by encoding information into an arbitrarily large block size, K. Yet, this is not attractive since it leads to a decoding delay that diverges to infinity. In this work, we address the question of whether RLNC can achieve throughput arbitrarily close to the capacity while yielding acceptable decoding delay. The main contributions of this work are:

- We find that the broadcast throughput of RLNC must vanish for any fixed K as N tends to infinity. We expose the cause of this behavior through a key example (see Section II), which motivates our search of a proper scaling of the block size K with increasing N.
- We introduce a tractable approximation of the original system by applying central limit theorem to properly scaled system parameters. We then establish tight bounds on the throughput and decoding-delay performance of the approximate system by using a combination of parametric bounding and optimization, as well as uniform bounds on the order statistics of Gaussian random variables.
- Our analysis reveals a *phase transition* in the performance of the approximate system that occurs at the block length scaling rate of $K = \Theta(\ln(N))$ with respect to the network size. Specifically, we show that if K increases slower than

[1] This asymptotic is with respect to increasing field size over which the data packets are defined (see Section III).

$\ln(N)$, then the broadcast throughput of RLNC converges to zero, and if K increases faster than $\ln(N)$, then the broadcast throughput of RLNC converges to the broadcast capacity of $(1-p)$.

- We provide extensive numerical results that compare the performance of the actual system performance to that of the approximate results. Our results uniformly show that the approximation is highly accurate, and the tight upper and lower bounds on the approximation hold for the actual system parameters.

These results collectively imply that RLNC can achieve throughput-delay tradeoff[2] of $((1-p), \Omega(\ln(N)))$. This is an attractive result as it indicates that as long as the coding block size scales super-logarithmically (i.e., very slowly) with the network size, the broadcast capacity is achievable with a simple policy such as RLNC.

The rest of the paper is organized as follows. In Section II, we overview some of the relevant work in this context and provide an example that motivates this work. After introducing the main system components in Section III, we provide our throughput and delay analysis of RLNC in Section IV. Our findings are confirmed in Section V through extensive numerical studies. Finally, our conclusions and remarks on future work are provided in Section VI.

II. Related Work and Motivating Example

Our model is similar to that considered in [5], [6], [7]. In [5], Ghaderi et al. quantify the reliability gain of RLNC for a fixed coding window size and show that this scheme significantly reduces the number of retransmissions in lossy networks compared to an end-to-end ARQ scheme. The delay performance gains of RLNC were observed by Eryilmaz et al. in [6]. They show that, for a fixed coding window size K, the network coding capability can lead to arbitrarily better delay performance as the system parameters (number of receivers) scale when compared to traditional transmission strategies without coding.

Also, in a similar setup as in this paper, it has been shown recently in [7] that there exists a phase transition with respect to decoding delay such that there exists a threshold on the number of transmissions below which the probability that a block of coded packets can be recovered by all the nodes in the network is close to zero. On the other hand, if the number of transmissions is slightly greater than the threshold, then the probability that every node in the network is able to reconstruct the block quickly approaches one.

All of the aforementioned works [5], [6], [7] study the gains of the network coding as the system size grows while the coding window size is held constant. In particular, they show that the decoding delay of RLNC scales as $O(\ln(N))$ for a fixed coding window size as $N \to \infty$. *However it can be seen that when the coding window size is held constant, the throughput of the system goes to zero as the system becomes large because each receiver gets a block of K packets in $O(\ln(N))$ time slots*. Therefore, it is important to study the system when K is scaled as a function of N.

The following example further highlights this point and motivates our investigation of the throughput-decoding delay tradeoff of RLNC:

Example 1: Consider a single source broadcasting blocks of K packets to N receivers in a rateless transmission. Each packet is a vector of length m over a finite field $\mathbb{F}_d$. In each time slot, the source broadcasts a random linear combination of K packets. Using random linear coding arguments introduced by Ho et al. [8], for a large enough field size d, it is sufficient for the receivers to receive approximately K coded packets to be able to decode the block.

Let the random variable $M[t]$ represent the number of receivers that have successfully decoded K packets in $t \geq K$ time slots. Let $r[t]$ represent the probability that any given receiver receives at least K packets in $t \geq K$ time slots. Then, $M[t]$ is a binomial random variable with probability of success $r[t]$, where $r[t] = \sum_{l=K}^{t} \binom{t}{l}(1-p)^l p^{t-l}$. Then $E(M[t]) = Nr[t]$. Here $r[t]$ represents the fraction of receivers that have successfully decoded K packets by the time t.

To compare the behavior of $r[t]$ as a function of t for different values of K, we define a normalized time variable, $s = \frac{t-K}{K}$. Accordingly, we define $r'[s] = r[Ks+K]$, which can be interpreted as the fraction of receivers that have successfully decoded a single packet in a block of K packets by s time slots. The comparison of $r'[s]$ for different K allows us to see, in a normalized time scale, the fraction of receivers that can decode an equivalent of a single packet from a batch of K.

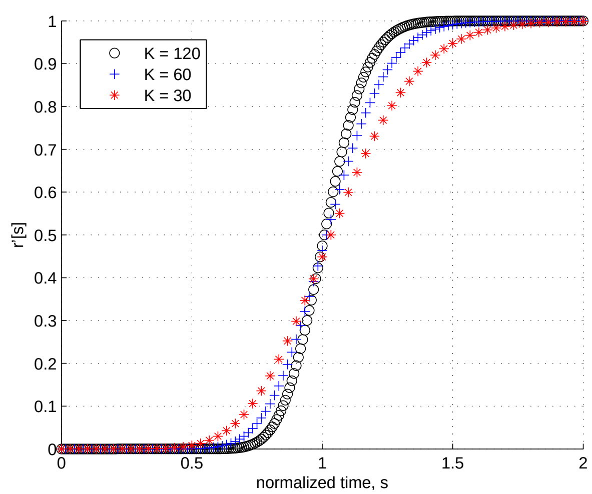

Fig. 1. Fraction of receivers that have successfully decoded a single packet in a block of K packets in s time slots, $r'[s]$ as a function of s for $p = 1/2$

We numerically evaluate $r'[s]$ as a function of s for different values of K as shown in the Figure 1 for the case where $p = 0.5$. It can be seen from the graph that for $K = 30$, a large

[2]We use the standard order notation: $g(n) = \omega(f(n))$ implies $\lim_{n\to\infty}(g(n)/f(n)) = \infty$; and $g(n) = \Omega(f(n))$ implies $\lim_{n\to\infty}(g(n)/f(n)) \geq c$ for some constant c; and $g(n) = \Theta(f(n))$ implies $\lim_{n\to\infty}(g(n)/f(n)) = c$ for some constant c.

fraction of users are served within a short duration and then the source takes a relatively longer time to serve the remaining small fraction of users towards the end of the transmission of the current block of K packets. On increasing K to $K = 60$, the graph becomes sharper indicating that the source serves a larger fraction of users in a shorter duration and takes lesser time to serve a smaller fraction of users towards the end of the transmission.

Ideally, we would like all the users to complete decoding together for an increase in throughput. This can be achieved by increasing K indefinitely as observed from Figure 1. However, this causes the decoding delay to increase indefinitely as well. Hence, it is important to understand the throughput-delay tradeoff as K scales as a function of N.

III. System Model

In this work, we study the basic wireless broadcast scenario depicted in Figure 2 that models the characteristics of cellular or satellite systems and serves as the fundamental building block for more general networks.

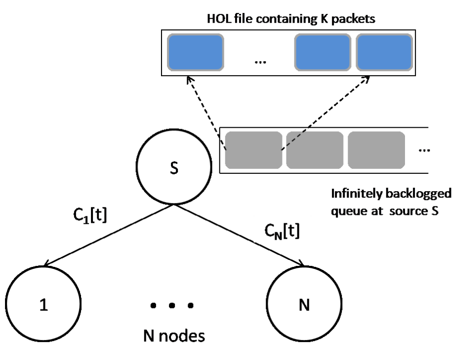

Fig. 2. A single source broadcasting to N receivers over erasure channels with probability of erasure, p in each time slot.

In particular, we consider a single source node, S, broadcasting an infinite backlog of data to N receivers over independent time-varying erasure channels. The data is encapsulated into *packets*, each represented as a vector of length m over a finite field $\mathbb{F}_d$. We assume a time-slotted operation of the system with $C_i[t] \in \{0, 1\}$ denoting the state of user is channel in slot t. We model $C_i[t]$ as a Bernoulli random variable with p being the probability that $C_i[t] = 0$ in any given time slot t. For simplicity, we assume that all channels are independent and identically distributed in our analysis for ease of exposition. A single packet may be broadcast in each time slot by the source and the transmission to the i^{th} user is successful only if $C_i[t] = 1$.

We consider the class of block coding strategies employed by the source, where data is transferred in blocks of K packets. Specifically, the source can start transmitting the next block only if the previous block is successfully transferred to all N receivers. Moreover, we focus on the Random Linear Network Coding (RLNC) strategy that is defined next.

Definition 1 (Random Linear Network Coding (RLNC)): In each time slot, the source transmits a random linear combination of the K packets in the Head-of-line (HOL) coding block (see Figure 2). In what follows, we refer to K as the *coding window (or block) size* of RLNC.

Using random linear coding arguments introduced by Ho et al. [8], for a large enough field size d, it is sufficient for the receivers to receive approximately K coded packets to be able to decode the block. It has been shown in [9] that random linear network coding is capacity achieving for multicast connections in an unreliable network setting as long as packets received on a link arrive according to a process that has an average rate. That is, for K sufficiently large, under the coding scheme defined in Definition 1, the (broadcast) capacity of our system is $(1 - p)$.

Next, we define the two metrics of interest in our analysis, namely throughput and decoding-delay.

Definition 2 (Decoding-Delay): We let $Y_i(K)$ denote the number of time slots it takes for the i^{th} receiver to decode a block of K packets under the RLNC scheme. Under the erasure channel model, it is easy to see that $Y_i(K)$ is a negative binomial random variables of order K. Then, the decoding delay for a given N and K under the RLNC scheme, denoted as $Z(N, K)$, is the time required to transmit all packets of the head-of-line (HOL) block to all the receivers. Hence, we have

$$Z(N, K) = \max_{1 \leq i \leq N} Y_i(K), \tag{1}$$

Definition 3 ((Broadcast) Throughput): We denote the number of packets transmitted by the source in a total of t slots by $R(t)$. Then, the (broadcast) throughput for a given N and K under RLNC scheme, denoted as $E[R(N, K)]$, is the long-term average number of successfully transferred data packets to all N receivers. Hence, from ergodicity, we have

$$E[R(N, K)] = \lim_{t \to \infty} \frac{R(t)}{t}. \tag{2}$$

The block transmission structure together with the independence of channel states across time allows us to model the RLNC operation as a renewal process with renewals at the start of each coding block formation. Hence, by defining a constant reward of K acquired in each renewal interval, we can utilize the main result from renewal theory [10] to write:

$$E[R(N, K)] = \lim_{t \to \infty} \frac{R(t)}{t} = \frac{K}{E[Z(N, K)]} \tag{3}$$

It is known that (e.g. [6]) the exact expression for $E[Z(N, K)]$ is as follows,

$$E[Z(N, K)] = K + \sum_{t=K}^{t=\infty} \left[1 - \left(\sum_{\tau=K}^{\tau=t} \binom{\tau - 1}{K - 1} p^{(\tau - K)} q^K \right)^N \right],$$

where $\binom{n}{m}$ gives the number of size m combinations of n elements and $q \triangleq (1 - p)$.

It can be seen from the above expression that, when K is a constant independent of N, the mean decoding delay $E[Z(N, K)]$ increases with N. Thus, for any fixed K, $E[R(N, K)]$ in (3) goes to zero as N approaches ∞.

However, the exact expression for the mean decoding delay is difficult to simplify further and does not provide any additional insight. This motivates us to study an approximation of the system performance that is later observed to be accurate.

IV. Throughput and Delay Estimation of RLNC

In this section, we define an approximation to the system of Section III and derive upper and lower bounds on the decoding delay and throughput of this new system. This enables us to understand the scaling of the coding window size with the number of receivers to guarantee a non vanishing throughput as the original system becomes large.

Let $Y(K)$ be a negative binomial random variable with mean $\mu(K) = \frac{K}{1-p}$ and variance $\sigma^2(K) = K\frac{p}{(1-p)^2}$. Define $\tilde{Y}(K) = (Y(K) - \mu(K))/\sigma(K)$. It is well known that if $Y(K)$ is a negative binomial random variable of order K and success probability $(1-p)$, then

$$Y(K) = \sum_{i=1}^{K} X_i, \tag{4}$$

where $X_i, i = 1, 2, \ldots, K$ is a sequence of independent geometric random variables with success probability $1-p$.

Hence one can invoke the central limit theorem for i.i.d. sequences which states that $\tilde{Y}(K)$ converges weakly to χ as $K \to \infty$ where χ has the standard normal distribution.

Noting the above convergence in distribution, we define a new system by replacing the negative binomial random variables with normal random variables in the original system. We expect that decoding delay and throughput of the new system would be a close approximation to that of RLNC, as will be confirmed in Section V.

Definition 4 (Approximate decoding-delay and throughput): We define *approximate decoding delay* $\tilde{Z}(N, K)$ and *approximate throughput* $E[\tilde{R}(N, K)]$ as follows,

$$\begin{aligned} \tilde{Z}(N,K) &= \frac{K}{1-p} + \frac{\sqrt{Kp}}{1-p} \max_{1\le i\le N} \tilde{\chi}_i, && (5)\\ E[\tilde{R}(N,K)] &= \frac{K}{E[\tilde{Z}(N,K)]}, && (6) \end{aligned}$$

where $\tilde{\chi}_i, i = 1, 2, \ldots, N$ are independent standard normal random variables.

Then, the main result of this paper is provided next:

Theorem 1: The approximate throughput $E[\tilde{R}(N, K)]$ shows the following behavior:

1) When $K = o(\ln(N))$, then $E[\tilde{R}(N, K)]$ goes to zero as N approaches ∞.
2) When $K = \Theta(\ln(N))$, then $E[\tilde{R}(N, K)]$ approaches a constant fraction of $(1-p)$ as N approaches ∞.
3) When $K = \omega(\ln(N))$, then $E[\tilde{R}(N, K)]$ approaches $(1-p)$ as N approaches ∞.

In order to prove the above theorem, we establish upper and lower bounds on the $E[\max_{1\le i\le N} \tilde{\chi}_i]$ in the following lemmas:

Lemma 1: For any $t > 0$ and standard normal random variables $\tilde{\chi}_i, i = 1, 2, \ldots, N$, we have that

$$E[\max_{1\le i\le N} \tilde{\chi}_i] \le \frac{1}{t}\left(\ln(N) + \frac{t^2}{2}\right) \tag{7}$$

The right-hand-side of inequality (7) can be minimized to get a better lower bound on $E[\max_{1\le i\le N} \tilde{\chi}_i]$ given by

$$E[\max_{1\le i\le N} \tilde{\chi}_i] \le \sqrt{2\ln(N)} \tag{8}$$

and the minimizing value of t, denoted by t^*, is given by,

$$t^* = \sqrt{2\ln(N)} \tag{9}$$

Proof: For any $y_1, \ldots, y_N$ and $t > 0$,

$$\begin{aligned} \max(y_1,\ldots,y_N) &= \frac{1}{t}(\ln(\max(e^{ty_1},\ldots,e^{ty_N}))) \\ &\le \frac{1}{t}(\ln(e^{ty_1}+\ldots+e^{ty_N})) \end{aligned}$$

Therefore, for standard normal random variables $\tilde{Y}_1, \ldots, \tilde{Y}_N$,

$$\begin{aligned} E[\max_{1\le i\le N} \tilde{\chi}_i] &\le \frac{1}{t}E[\ln(e^{t\tilde{\chi}_1}+\ldots+e^{t\tilde{\chi}_N})] \\ &\le \frac{1}{t}(\ln(E[e^{t\tilde{\chi}_1}+\ldots+e^{t\tilde{\chi}_N}])) \\ &= \frac{1}{t}(\ln(N)+\ln(E[e^{t\tilde{\chi}_1}])) \\ &= \frac{1}{t}\left(\ln(N)+\frac{t^2}{2}\right) \end{aligned}$$

■

Lemma 1 is similar to the Lemma 2 in [**?**] which states that given $X_1, \ldots, X_N$ i.i.d random variables, for any $t > 0$, $E[maxX_1, \ldots, X_N] \le \frac{1}{t}(\ln(N) + \ln(E[e^{tX_1}]))$. We use this inequality in the case of standard normal random variables and optimize the upper bound by minimizing with respect to t noting that the moment generating function of the standard normal random variables converges for all $t > 0$.

Lemma 2: Let $\tilde{\chi}_i, i = 1, 2, \ldots, N$ be a sequence of i.i.d. standard normal random variables. There exists a universal constant $0 < C \le \sqrt{2}$ such that

$$E[\max_{1\le i\le N} \tilde{\chi}_i] \ge C\sqrt{\ln(N)} \tag{10}$$

Proof: The proof of (10) follows from [11], Lemma (4.10). The fact that $C \le \sqrt{2}$ follows from the upper bound from Lemma 1 that $C \le \sqrt{2}$

■

Proof: (Proof of the theorem) From Lemma 1 and Lemma 2, we have that, for some $0 < C \le \sqrt{2}$:

$$\begin{aligned} \frac{K}{1-p} + \frac{\sqrt{CKp\ln(N)}}{1-p} &\le E[\tilde{Z}(N,K)] && (11)\\ &\le \frac{K}{1-p} + \frac{\sqrt{2Kp\ln(N)}}{1-p} \end{aligned}$$

And,

$$(1-p)\left[1+\frac{\sqrt{2p\ln(N)}}{\sqrt{K}(1-p)}\right]^{-1} \leq E[\tilde{R}(N,K)] \qquad (12)$$
$$\leq (1-p)\left[1+\frac{\sqrt{Cp\ln(N)}}{\sqrt{K}(1-p)}\right]^{-1}$$

Theorem 1 then follows from the definition of the approximate throughput and the above inequalities.

■

V. Numerical Results

In this section, we provide extensive numerical results to compare the behavior of the approximate system studied in Section IV to that of the RLNC in an actual system. Our results uniformly suggest the accuracy of the approximation and the applicability of the upper and lower bounds of the approximation to the actual system behavior.

As a representative setup, we let the OFF probability of erasure channels p to be 0.1. We note that the scaling behavior of the throughput and decoding-delay do not change for any other choice of p. Also, note that the broadcast capacity for this choice of p is $(1-p) = 0.9$. Our numerical results are presented under two different scenarios, the first focusing on confirming the phase transition of the throughput scaling, and the second focusing on confirming the applicability of lower and upper bounds obtained in Theorem 1.

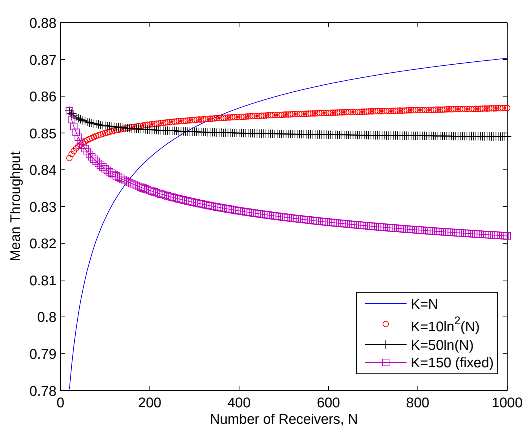

Fig. 3. Throughput behavior under different scalings of K with N.

Study 1) Phase transition: In this study, we explore the phase transition law that is suggested by Theorem 1. To that end, Figure 3 depicts the mean (broadcast) throughput of RLNC in the actual system operation with increasing N for different types of scaling of K. We see that this result is in perfect agreement with the phase transition law: when $K = 150$ and therefore scales slower than $\ln(N)$, we see that the throughput decays towards zero; when $K = 50\ln(N)$, i.e. $K = \Theta(\ln(N))$, the throughput converges to a constant level as suggested by the approximate analysis; when $K = 10\ln^2(N)$ or N, i.e. $K = \omega(\ln(N))$, the throughput increases toward the broadcast capacity. The latter two results also reveal that the converge rate the performance to the capacity may be increased by selecting a faster scaling of K with respect to N.

Thus, Study 1 confirms the phase transition law suggested by the approximate analysis. The next study is aimed at studying the accuracy of the lower and upper bounds in the proof of Theorem 1 (c.f. (11) and (12)) for different scalings of K with N.

Study 2) Lower and upper bounds: Here, we consider two different scalings of K with respect to N, and compare both the throughput and the decoding-delay performance of the actual system behavior to the system performance under the Gaussian approximation, and to the lower and upper bounds (11) and (12) obtained for the approximate system. In particular, we study the cases when $K = 50\ln(N)$ and $K = N$, which are already demonstrated in Figure 3 to converge to a constant and to 0.9, respectively.

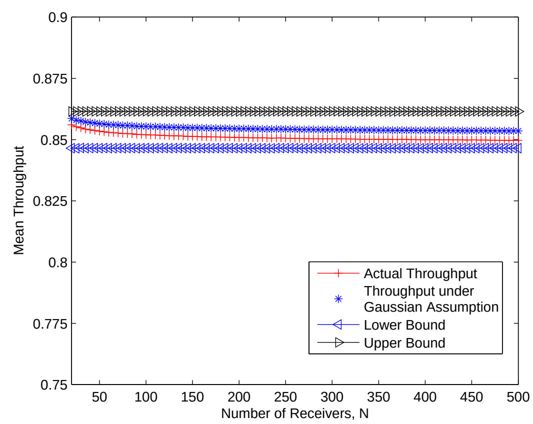

Fig. 4. Comparing Actual Throughput to Obtained Upper and Lower Bounds for $p = 0.1$ and $K = 50\ln(N)$

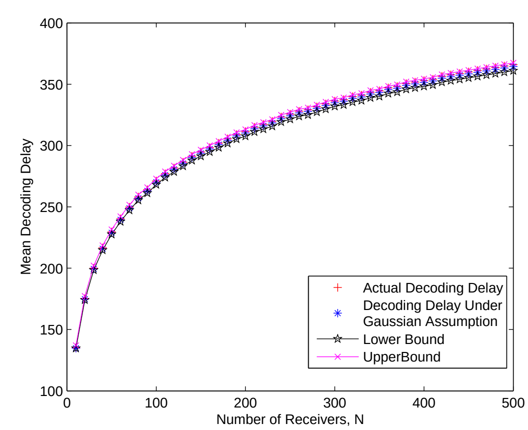

Fig. 5. Comparing Actual Mean Completion Time to Obtained Upper and Lower Bounds for $p = 0.1$ and $K = 50\ln(N)$

Figures 4 and 5 respectively depict the throughput and decoding-delay performance with $K = 50\ln(N)$ of the actual system behavior together with the approximation and the lower and upper bounds. These demonstrate both the tightness the lower and upper bounds and the fact that the actual system performance is also bounded by them. We also see that a throughput of approximately 0.85 is achievable with this scaling, leading to a decoding delay that scales only logarithmically with the network size.

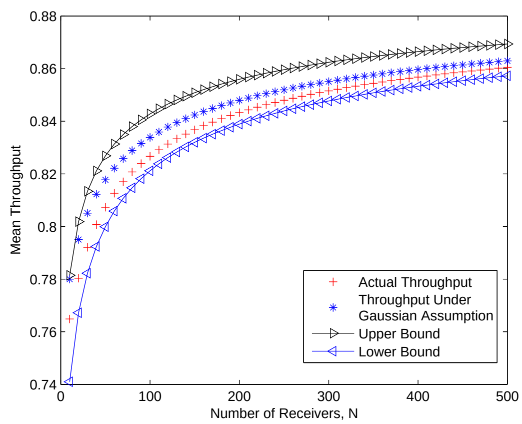

Fig. 6. Comparing Actual Throughput to Obtained Upper and Lower Bounds for $p = 0.1$ and $K = N$

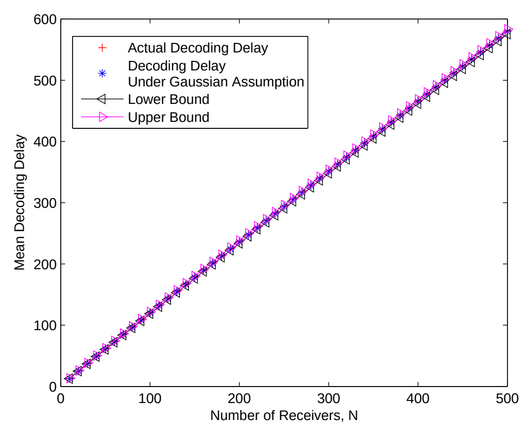

Fig. 7. Comparing Actual Mean Completion Time to Obtained Upper and Lower Bounds for $p = 0.1$ and $K = N$

In comparison, Figures 6 and 7, respectively, depict the throughput and decoding-delay performance of the actual and approximate systems when $K = N$ along with the bounds. Again, we observe that the bounds are tight and are applicable to the actual system performance, as predicted. In this fast scaling scenario, we also observe that the throughput increases towards the capacity of 0.9 instead of converging to a constant level as in Figure 4. Yet, this asymptotic optimality occurs at the cost of linearly increasing decoding-delay performance.

Overall, these numerical studies collectively confirm the accuracy of estimating the RLNC performance using the Gaussian approximations. The rigorous proof of this result is part of our future research. This connection to a tractable formulation is expected to be of paramount importance in the further analysis and the extension of these results to more involved setups.

VI. Conclusions

We have investigated the throughput and decoding delay performance of RLNC in a wireless broadcast setting as the coding window size K scales as N. We noted that the broadcast throughput of RLNC vanishes for any fixed K as the system size increases. Hence, it is important to understand the scaling of K as a function of N that will guarantee a non-vanishing throughput. To this end, we defined an approximate system by appealing to the central limit theorem and replacing the negative binomial random variables in the analysis of decoding delay with normal random variables.

Our analysis revealed a phase transition in the performance of the approximate system, namely, if K increases slower than $\ln(N)$, the throughput goes to zero as N increases. However, on increasing K faster than $\ln(N)$, the throughput approaches the maximum achievable broadcast throughput of (1-p). Also, $K = \Theta(\ln(N))$ ensures a constant fraction of the maximum achievable broadcast throughput for the approximate system.

We have shown through numerical results that our analysis closely approximates the original system and in fact, the performance of the original system exhibits a phase transition indicating that our results apply to the original system as well.

References

[1] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 36, pp. 1936–1948, December 1992.

[2] X. Lin and N. Shroff, "The impact of imperfect scheduling on cross-layer rate control in multihop wireless networks," in *Proceedings of IEEE Infocom*, Miami, FL, March 2005.

[3] A. Eryilmaz and R. Srikant, "Joint congestion control, routing and mac for stability and fairness in wireless networks," *IEEE Journal on Selected Areas in Communications, special issue on Nonlinear Optimization of Communication Systems*, vol. 14, pp. 1514–1524, August 2006.

[4] A. Eryilmaz, A. E. Ozdaglar, M. Médard, and E. Ahmed, "On the delay and throughput gains of coding in unreliable networks," *IEEE Transactions on Information Theory*, vol. 54, no. 12, pp. 5511–5524, 2008.

[5] M. Ghaderi, D. F. Towsley, and J. F. Kurose, "Reliability gain of network coding in lossy wireless networks," in *INFOCOM*, 2008, pp. 2171–2179.

[6] A. Eryilmaz, A. E. Ozdaglar, M. Médard, and E. Ahmed, "On the delay and throughput gains of coding in unreliable networks," *IEEE Transactions on Information Theory*, vol. 54, no. 12, pp. 5511–5524, 2008.

[7] W. Xiao and D. Starobinski, "Extreme value fec for reliable broadcasting in wireless networks," in *INFOCOM*, 2009.

[8] T. Ho, "Networking from a network coding perspective," Ph.D. dissertation, MIT, 2004.

[9] D. S. Lun, M. Médard, R. Koetter, and M. Effros, "On coding for reliable communication over packet networks," *CoRR*, vol. abs/cs/0510070, 2005.

[10] S. Ross, *Stochastic Processes*. John Wiley and Sons, 1983.

[11] Adler, *An Introduction to Continuity, Extrema, and Related Topics for General Gaussian Processes*. Institute of Mathematical Statistics, 1990.