Graham, M., Ganesh, A., & Piechocki, R. (2019). Fountain Coding Enabled Data Dissemination for Connected and Automated Vehicles. In *2019 IEEE 89th Vehicular Technology Conference, VTC Spring 2019 - Proceedings* [8746545] (IEEE Vehicular Technology Conference; Vol. 2019-April). Institute of Electrical and Electronics Engineers (IEEE). <https://doi.org/10.1109/VTCSpring.2019.8746545>

Peer reviewed version

Link to published version (if available):
[10.1109/VTCSpring.2019.8746545](https://doi.org/10.1109/VTCSpring.2019.8746545)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via IEEE at <https://ieeexplore.ieee.org/document/8746545>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Fountain Coding Enabled Data Dissemination for Connected and Automated Vehicles

Mark A. Graham
School of Mathematics
University of Bristol
Bristol, United Kingdom
mark.graham@bristol.ac.uk

Ayalvadi Ganesh
School of Mathematics
University of Bristol
Bristol, United Kingdom
a.ganesh@bristol.ac.uk

Robert J. Piechocki
Faculty of Engineering
University of Bristol
Bristol, United Kingdom
r.j.piechocki@bristol.ac.uk

Abstract—The emergence of new connectivity services for automated transportation marks a paradigm shift for the operation of wireless networks. Furthermore, the advent of blockchain technology promises to enable a plethora of smart mobility services, which are not contingent on any central authorities. Concepts such as distributed ledger require efficient and reliable data dissemination between vehicles. Traditional techniques based on Automatic Repeat Request (ARQ) are well known to scale poorly in all-cast networks due to the feedback implosion problem. Fountain and network coding techniques are arguably the most promising alternative solutions. In this paper we derive new analytical bounds on transmit message lengths and quantify bandwidth delay trade-offs for fountain coding based data dissemination for CAVs.

I. INTRODUCTION

Connected and Automated Vehicles (CAV) are poised to have a range of positive impacts of the future of mobility: from reducing journey times and CO2 emissions to increasing the road safety. But realising this vision will require reliable and low-latency communication of sensor, telemetry, forensic and other types of data between vehicles. Channel conditions in Vehicle to Vehicle (V2V) and Vehicle to Infrastructure (V2I) networks are very challenging, with high packet loss rates [1] due to contention or other vehicles obstructing the communication path [2]. The rapid movement of all vehicles involved also causes frequent changes in network topology [3]. The objective is to ensure reliable dissemination of data transmitted by each vehicle to all other vehicles within some specified range or neighborhood.

In more detail, we consider a single source broadcasting an infinite stream of packets to multiple receivers, over independent broadcast channels. We seek to characterise the range of achievable throughputs and delays, and practical schemes to achieve them. The feedback implosion problem suffered by ARQ is well known, benefit for systems with more than one user.

The ARQ method, in which erased packets are re-transmitted on the receipt of feedback, is well known to suffer a feedback implosion problem for broadcasts. In its basic form *Fountain coding* involves the transmission of a stream of coded linear combinations of message packets until

all receivers are able to decode all the message packets. A common approach is to draw the coefficients of the linear combination at random, as in *Random Linear Network Coding (RLNC)* [4]. Once a receiver has received as many linearly independent coded packets as packets in the message, the original message can be decoded by solving the resulting linear system. The coefficients may be chosen randomly in such a way as to produce sparse linear combinations, as in the Luby Transform (LT) code, which allows the message to be decoded in linear time using the belief propagation algorithm [5]. Fountain coding has the advantage of requiring no more feedback than a single acknowledgement (ack) when a receiver has acquired enough packets to enable it to decode the message.

The authors of [4] show that RLNC is throughput optimal for broadcast. The authors of [6] derive a bound on the number of transmissions required for all users to receive the message, with probability exceeding a specified threshold. They also present a computationally expensive algorithm for minimising the expected completion time by optimally choosing N_i , and a computationally feasible heuristic.

The analysis is extended to time-varying channel erasure probabilities in [7]. The authors derive an exact expression for the expected number of transmissions required for successful reception of a fixed length message by all broadcast receivers, when RLNC is employed. A comparison of numerical values of this expectation, as well as numerical and simulated expected completion times for a number of other scheduling algorithms detailed in the paper, shows that RLNC is superior in all cases. Paper [8] considers RLNC applied to a time invariant broadcast erasure channel, where the transmitter is linked to each receiver by an independent erasure channel, with distinct erasure probabilities. The author derives a lower bound for the number of transmissions required for successful reception of a fixed length message by all broadcast receivers, and shows this to be close to the true expressions for practical system parameters in simulations. The result however is cumbersome, and a simpler, more insightful approximation is derived by making additional assumptions.

In this paper we compare the performance of fountain coding with ARQ for broadcast erasure channels. For a message of fixed length, we obtain analytical bounds on the

number of packet transmissions until all receivers can decode the message, both with ARQ and with fountain coding. For channels with identical erasure probabilities, these bounds can be simplified in the joint asymptotic regime in which both the message length and the number of receivers tend to infinity, yielding simple and insightful expressions. The analytical results are complemented by extensive simulations.

II. SYSTEM MODEL

We consider a broadcast model where vehicles attempt to transmit messages to all other vehicles. Each transmitter seeks to communicate the same message, consisting of a string of packets, to each of n receivers. One packet can be transmitted over the broadcast channel in each time step; each receiver either receives it fully and free of error, or receives nothing. The erasures are assumed to be mutually independent across receivers and across time slots; q_i denotes the erasure probability at receiver i . The same channel model is considered in [7], [9], [10], [11], and the same questions addressed.

In ARQ, each packet is broadcast repeatedly, until all receivers have ACK'ed its receipt. We assume ACKs are received without error, and incur no overhead. We also make idealised assumptions about fountain coding, namely that it is perfect and that every transmitted packet is linearly independent of (relevant) previous transmissions. But if the field size is large, and the number of coded packets required is not too large, then this probability is close to one; see [7], [4] for an analysis that takes this into account.

The maximum rate achievable by any coding scheme is $1/(1 - \max_{i=1}^n q_i)$, as that is the channel capacity to the worst-off receiver. It is well known that RLNC can achieve this rate; see [4], for instance. In contrast, ARQ requires $X_i \sim \text{Geom}(1 - q_i)$ transmissions of a packet until receiver i gets it, and $\max_{i=1}^n X_i$ transmissions for all receivers to get it. This number typically scales as $\log n$ (if $\min_{i=1}^n q_i$ and $\max_{i=1}^n q_i$ are bounded away from zero and one, uniformly in n), and so the throughput of ARQ scales as $1/\log n$, which vanishes as the number of receivers increases to infinity. It is less obvious how the delay of these schemes scale with k and n , which is the topic of the next section.

III. THROUGHPUT DELAY TRADE-OFFS

We assume henceforth that the erasure probabilities q_i are identical across receivers, and denote it by q . While this is not essential, it leads to simple expressions that yield insight on the scaling laws relating throughput and delay.

We begin by recalling a standard result about large deviations of binomial random variables, which is an immediate consequence of Sanov's theorem.

Lemma 1. *Suppose that X is a binomially distributed random variable with parameters (n, p) , which we denote by $X \sim \text{Bin}(n, p)$. Then,*

$$\begin{cases} \mathbb{P}(X > nq) \leq \exp(-nH(q; p)), & \forall q > p, \\ \mathbb{P}(X < nq) \leq \exp(-nH(q; p)), & \forall q < p, \end{cases} \quad (1)$$

where

$$H(\beta; \alpha) = \beta \log \frac{\beta}{\alpha} + (1 - \beta) \log \frac{1 - \beta}{1 - \alpha}$$

denotes the relative entropy or Kullback-Leibler (KL) divergence of the Bernoulli(β) distribution with respect to the Bernoulli(α) distribution.

The main result in this section is the following theorem, which shows that fountain coding can achieve the full range of possible rates, from zero up to channel capacity, with latency that is logarithmic in the number of receivers (albeit with a constant prefactor that depends on the rate).

Theorem 1. *Consider a sequence of systems where the n^{th} system has a single transmitter communicating with n receivers over a broadcast erasure channel with constant erasure probability $q \in (0, 1)$. Erasures are independent across receivers and time slots. Suppose the transmitter employs perfect fountain coding over blocks of k_n packets. Denote by $T = T(n, k_n)$ the random number of time slots until all n receivers have decoded all k_n packets in the message. We have the following:*

$$\text{If } \frac{k_n}{\log n} \rightarrow \alpha \geq 0, \text{ then } \frac{T(n, k_n)}{\log n} \xrightarrow{P} \beta_\alpha,$$

where $\beta_\alpha := \inf \left\{ \beta > \frac{\alpha}{1-q} : \beta H\left(\frac{\alpha}{\beta}; 1-q\right) > 1 \right\}$ and $\xrightarrow{P}$ denotes convergence in probability. Moreover, the function $\beta \mapsto \beta H\left(\frac{\alpha}{\beta}; 1-q\right)$ is non-decreasing on $\left[\frac{\alpha}{1-q}, \infty\right)$.

It is easy to see from properties of the relative entropy function that the set over which the infimum in the definition of β_α is taken is non-empty for all $\alpha \geq 0$, and that the infimum is attained; consequently, β_α is finite.

Proof. By the assumption of perfect fountain coding, receiver i has decoded the message by time ℓ if it has received at least k_n packets by this time, i.e., suffered no more than $\ell - k_n$ erasures. The number of erasures in ℓ time slots is binomially distributed with parameters (ℓ, q) . Hence, letting T_i denote the random time at which receiver i decodes the message, we have $\mathbb{P}(T_i > \ell) = \mathbb{P}(\text{Bin}(\ell, 1 - q) < k_n)$. As

$$T(n, k_n) = \max_{i=1}^n T_i,$$

it follows from the union bound that

$$\mathbb{P}(T(n, k_n) > \ell) \leq n \mathbb{P}(\text{Bin}(\ell, 1 - q) < k_n).$$

Hence, by Lemma 1, we have for $\ell > \frac{k_n}{1-q}$ that

$$\log \mathbb{P}(T(n, k_n) > \ell) \leq \log n - \ell H\left(\frac{k_n}{\ell}; 1 - q\right). \quad (2)$$

It is well-known, and easily verified, that

$$H(q; p) = \sup_{\theta \in \mathbb{R}} \theta q - \log(pe^\theta + 1 - p),$$

and, moreover, that the supremum is attained at $\theta \leq 0$ for $q \leq p$ and at $\theta \geq 0$ for $q \geq p$. Hence, we have for $\beta \geq \frac{\alpha}{1-q}$ that

$$\begin{aligned}\beta H\left(\frac{\alpha}{\beta}; 1-q\right) &= \beta \sup_{\theta \leq 0} \frac{\theta \alpha}{\beta} - \log((1-q)e^\theta + q) \\ &= \sup_{\theta \leq 0} \theta \alpha - \beta \log((1-q)e^\theta + q).\end{aligned}$$

Now, for $\theta \leq 0$, $\log((1-q)e^\theta + q) \leq 0$, and so the last expression on the right above is a supremum of non-decreasing functions of β ; hence, it is also a non-decreasing function of β . This proves the last claim of the theorem. Fix $\epsilon > 0$, $\beta > (1+\epsilon)\beta_\alpha$ and take $\ell_n = \lceil \beta \log n \rceil$, where $\lceil x \rceil$ denotes the smallest integer that is no smaller than x . By the assumption that $k_n/\log n$ tends to α , and the monotonicity of $\beta \mapsto \beta H(\frac{\alpha}{\beta}; 1-q)$ on $[\frac{\alpha}{1-q}, \infty)$ established above, we have for all n sufficiently large that

$$\begin{aligned}\ell_n H\left(\frac{k_n}{\ell_n}; 1-q\right) &\geq (1+\epsilon) \log n \beta_\alpha H\left(\frac{\alpha}{\beta_\alpha}; 1-q\right) \\ &\geq (1+\epsilon) \log n,\end{aligned}$$

where the last equality holds by the definition of β_α and the continuity of $H(\cdot; 1-q)$. Substituting the above inequality into (2), we get

$$\begin{aligned}\limsup_{n \rightarrow \infty} \mathbb{P}(T(n, k_n) > (1+\epsilon)\beta_\alpha \log n) \\ \leq \limsup_{n \rightarrow \infty} n^{-\epsilon} = 0.\end{aligned}\quad (3)$$

We need a corresponding lower bound on $T(n, k_n)$. Observe that T_i , the number of transmissions until receiver i gets k_n packets, satisfies

$$\begin{aligned}\mathbb{P}(T_i \leq \ell) &= \mathbb{P}(\text{Bin}(\ell, 1-q) \geq k_n) \\ &\leq 1 - \binom{\ell}{k_n} (1-q)^{k_n} q^{\ell-k_n}.\end{aligned}$$

Now, using Stirling's formula, we obtain after some routine simplifications that, for arbitrary $\epsilon > 0$ and all ℓ and k_n sufficiently large,

$$\mathbb{P}(T_i \leq \ell) \leq 1 - \frac{1-\epsilon}{\sqrt{2\pi\ell}} \exp\left(-\ell H\left(\frac{k_n}{\ell}; 1-q\right)\right). \quad (4)$$

Next, let $\ell_n, n \in \mathbb{N}$ be a sequence satisfying $\frac{\alpha}{1-q} \leq \frac{\ell_n}{\log n} \leq (1-\epsilon)\beta_\alpha$; such a sequence exists for ϵ sufficiently small. Since $k_n/\log n$ tends to α , observe that for all n sufficiently large, we have

$$\begin{aligned}\ell_n H\left(\frac{k_n}{\ell_n}; 1-q\right) &\leq (1-\epsilon) \log n \beta_\alpha H\left(\frac{\alpha}{\beta_\alpha}; 1-q\right) \\ &\leq (1-\epsilon) \log n.\end{aligned}$$

Substituting this into (4), we get for all n sufficiently large that

$$\mathbb{P}(T_i \leq (1-\epsilon)\beta_\alpha \log n) \leq 1 - \frac{(1-\epsilon)^2}{\sqrt{2\pi\beta_\alpha \log n}} n^{-(1-\epsilon)}. \quad (5)$$

The total number of transmissions until all n receivers can decode the k_n message packets is given by $T(n, k_n) =$

$\max_{i=1}^n T_i$. Moreover, the random variables T_i are mutually independent by the assumption that erasures on channels to distinct receivers are mutually independent. Hence,

$$\mathbb{P}(T(n, k_n) \leq \ell) = \mathbb{P}(\forall i : T_i \leq \ell) = \mathbb{P}(T_i \leq \ell)^n.$$

Substituting (5) into the above, we obtain after some simple manipulations that

$$\limsup_{n \rightarrow \infty} \mathbb{P}(T(n, k_n) < (1-\epsilon)\beta_\alpha \log n) = 0. \quad (6)$$

In conjunction with (3), this completes the proof of the theorem. $\square$

Remarks. The above theorem describes the achievable trade-off between throughput and delay. At one extreme, if we take $k_n = 1$ for all n , then the message consists of a single packet, and fountain coding reduces to ARQ. This corresponds to $\alpha = 0$ in the theorem statement. The theorem tells us that

$$\beta_0 = \frac{1}{H(0; 1-q)} = \frac{-1}{\log q},$$

i.e., it needs $-\log n / \log q$ packet transmissions to recover the message. This is the minimum achievable latency, and is achieved for throughput vanishing in the limit as the number of receivers, n , tends to infinity. The theorem further tells us that throughputs arbitrarily close to capacity are achievable while keeping latencies of the same order, namely logarithmic in n . In particular, if we take $k_n \sim \alpha \log n$, then we incur a latency of $\beta_\alpha \log n$ while achieving a throughput of α/β_α . As α increases to infinity, so does β_α , while the ratio α/β_α tends to $1-q$, which is the channel capacity. In other words, as throughput approaches capacity, the latency occurred becomes an arbitrarily large multiple of $\log n$.

The relationship between the delay (scaled by $\log n$) and throughput is plotted in Figure 1, for three different values of the erasure probability q . The figure shows that higher throughputs incur higher delays, and that the delay blows up as throughput approaches capacity. Moreover, delay increases with the erasure probability.

We conclude this section with a heuristic calculation of $T(n, k_n)$ based on the Central Limit Theorem (CLT). Even though Theorem 1 gives a precise asymptotic expression for $T(n, k_n)$, it is not clear *a priori* how large n and k_n need to be to yield a good approximation. We now present an alternative approximation.

Let X_i denote a random variable with the distribution of the number of packet transmissions required for receiver i to receive a single packet. Then, $X_i \sim \text{Geom}(1-q)$, and so $\mathbb{E}[X_i] = 1/(1-q)$ and $\text{Var}(X_i) = q/(1-q)^2$. Now $T_i(k)$, defined as the time until receiver i obtains k distinct packets, is the sum of k iid copies of X_i . Hence, for large k , we have by the CLT that

$$\frac{1-q}{\sqrt{kq}} \left(T_i(k) - \frac{k}{1-q} \right) \Rightarrow Z \text{ as } k \rightarrow \infty,$$

where $\Rightarrow$ denotes convergence in distribution, and Z denotes a standard normal random variable. Let Φ denote the cdf of Z , i.e., $\Phi(x) = \mathbb{P}(Z \leq x)$.

Fig. 1. Delay vs. throughput of perfect fountain coding.

Fig. 2. Delay coefficients from Theorem 1 (β_α) and the CLT approximation in Equation. 9 (γ_α).

Now $T(n, k) = \max_{i=1}^n T_i(k)$, and the $T_i(k)$ are mutually independent. Hence, $\mathbb{P}(T(n, k) \leq x) = (\mathbb{P}(T(k) \leq x))^n$. Using the CLT to approximate the cdf of $T_i(k)$ as above, we get

$$\mathbb{P}\left(T(n, k) \leq \frac{k + \sqrt{kqx \log n}}{1 - q}\right) \approx \left(\Phi\left(\sqrt{x \log n}\right)\right)^n. \quad (7)$$

The reason that this is not a limit theorem is that the CLT establishes convergence of distribution in the bulk, whereas we are using it inappropriately to approximate the distribution in the tail.

Next, using the inequality

$$1 - \Phi(z) \leq \frac{1}{z\sqrt{2\pi}} \exp\left(-\frac{z^2}{2}\right),$$

which can be obtained by integrating the normal density by

parts, we have from (7) that

$$\begin{aligned} \mathbb{P}\left(T(n, k) \leq \frac{k + \sqrt{kqx \log n}}{1 - q}\right) \\ \approx \left(1 - \frac{n^{-x/2}}{\sqrt{2\pi x \log n}}\right)^n \approx \exp\left(-\frac{n^{1-(x/2)}}{\sqrt{2\pi x \log n}}\right). \end{aligned}$$

The term in the exponent tends to zero if $x > 2$ and to $-\infty$ if $x < 2$. This suggests the heuristic

$$T(n, k) \approx \frac{k + \sqrt{2kq \log n}}{1 - q}. \quad (8)$$

Comparing this result with Theorem 1, we see that if we fix $\alpha > 0$ and take $k_n = \lceil \alpha \log n \rceil$, then (8) suggests

$$T(n, k_n) \approx \gamma_\alpha \log n, \text{ where } \gamma_\alpha = \frac{\alpha + \sqrt{2\alpha q}}{1 - q}, \quad (9)$$

while Theorem 1 yields the asymptotically correct expression $T(n, k_n) \approx \beta_\alpha \log n$. We have plotted both β_α and γ_α against α in Figure 2, which shows that they are very close to each other, but diverge as α increases. The figure leads us to conjecture that $\beta_\alpha \geq \gamma_\alpha$ for all $\alpha \geq 0$.

IV. SIMULATION RESULTS

In order to compare the performance of fountain coding with that of ARQ, broadcasts of messages consisting of varying numbers of packets to varying numbers of users over the broadcast erasure channel described in Section II were simulated using each of the techniques. The simulations were implemented in Python. Each model and parameter combination was simulated for 368640 messages. Simulations were carried out for three different erasure probabilities, $q = 0.1, 0.2$ and 0.3 , for $n = 20$ receivers, and for a wide range of message sizes, ranging from $k = 1$ to $k = 100$ packets. In each case, the total number of packet transmissions, $T(n, k)$, required until all receivers were able to decode the message was obtained from the simulations. The excess latency, defined as the amount by which $T(n, k)$ exceeded the expected minimum number of packets, $k/(1-q)$, required to transmit the message to a single receiver, was calculated. This is a measure of the overhead caused by having multiple receivers. We compare how this overhead differs between ARQ and fountain coding, and also how it depends on the message size.

Figure 3 compares the average excess latency of fountain coding and ARQ in simulations. Excess latency is defined as $T(n, k) - \frac{k}{1-q}$, and the average over 368640 simulations is computed. Notice how the excess latency of ARQ is much larger than that of fountain coding. Moreover, the excess latency of ARQ grows linearly with the message size k , whereas that of fountain coding grows sub-linearly, approximately as $\sqrt{k}$ as predicted by Equation. (9) in Section III.

Figure 4 plots the average excess latency observed in the simulations, for erasure probability $q = 0.3$ and $n = 20$ receivers, as a function of the message size k . The theoretical predictions for the same quantity from Theorem 1 and the CLT approximation, Equation. 9 are also shown on the same plot. Even though the theoretical results are asymptotic, in a

Fig. 3. Graphic comparing the average excess latency of fountain coding and ARQ in 368640 simulations, $n = 20$.

Fig. 4. Graphic comparing the average excess latency of fountain coding in 368640 simulations, $n = 20$, with theoretical predictions from Section III.

limiting regime in which k and n tend to infinity, the figure shows that they give quite good predictions even for rather small values of k and n .

V. CONCLUSIONS

We have shown that feedback based solutions such as ARQ scale poorly for broadcast channels for increasing numbers of receivers, even under the generous assumption of perfect error and contention free feedback channels, and have quantified the bandwidth delay trade-offs for fountain codes operating on the same channel. Using simulations, we have shown that our asymptotic expressions are useful as approximations with good accuracy for finite messages and finite numbers of receivers, making these results insightful for parameter selection in practical all-cast networks enabling data dissemination for future connected vehicles.

VI. ACKNOWLEDGEMENTS

The authors would like to thank Stephen Wales, Roke Manor Research, for helpful discussions.

REFERENCES

- [1] F. Bai and B. Krishnamachari, "Exploiting the wisdom of the crowd: localized, distributed information-centric vanets [topics in automotive networking]," *IEEE Communications Magazine*, vol. 48, no. 5, 2010.
- [2] M. Boban, T. T. Vinhoza, M. Ferreira, J. Barros, and O. K. Tonguz, "Impact of vehicles as obstacles in vehicular ad hoc networks," *IEEE journal on selected areas in communications*, vol. 29, no. 1, pp. 15–28, 2011.
- [3] N. Lu, N. Cheng, N. Zhang, X. Shen, and J. W. Mark, "Connected vehicles: Solutions and challenges," *IEEE internet of things journal*, vol. 1, no. 4, pp. 289–299, 2014.
- [4] D. S. Lun, M. Médard, R. Koetter, and M. Effros, "On coding for reliable communication over packet networks," *Physical Communication*, vol. 1, no. 1, pp. 3–20, 2008.
- [5] M. Luby, "Lt codes," in *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*. IEEE, 2002, pp. 271–280.
- [6] D. E. Lucani, M. Médard, and M. Stojanovic, "Broadcasting in time-division duplexing: A random linear network coding approach," in *Network Coding, Theory, and Applications, 2009. NetCod'09. Workshop on*. IEEE, 2009, pp. 62–67.
- [7] A. Eryilmaz, A. Ozdaglar, and M. Medard, "On delay performance gains from network coding," in *Information Sciences and Systems, 2006 40th Annual Conference on*. IEEE, 2006, pp. 864–870.
- [8] E. Tsimbaló, A. Tassi, and R. J. Piechocki, "Reliability of multicast under random linear network coding," *IEEE Transactions on Communications*, vol. 66, no. 6, pp. 2547–2559, 2018.
- [9] B. Shrader and A. Ephremides, "On the queueing delay of a multicast erasure channel," in *IEEE Information Theory Workshop, Chengdu, 2006*, pp. 423–427.
- [10] B. T. Swapna, A. Eryilmaz, and N. B. Shroff, "Throughput-delay analysis of random linear network coding for wireless broadcasting," in *IEEE International Symposium on Network Coding (NetCod)*, 2010.
- [11] N. Xie and S. Weber, "Delay on broadcast erasure channels under random linear combinations," *IEEE Transactions on Information Theory*, vol. 63, no. 3, pp. 1631–1661, 2017.