



On the Throughput and Energy Benefits of Network Coded Cooperation

Hernandez, Nestor; Heide, Janus; Roetter, Daniel Enrique Lucani; Fitzek, Frank

Published in:
Cloud Networking (CloudNet), 2014 IEEE 3rd International Conference on

DOI (link to publication from Publisher):
[10.1109/CloudNet.2014.6968982](https://doi.org/10.1109/CloudNet.2014.6968982)

Publication date:
2014

Document Version
Early version, also known as pre-print

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Hernandez, N., Heide, J., Roetter, D. E. L., & Fitzek, F. (2014). On the Throughput and Energy Benefits of Network Coded Cooperation. In *Cloud Networking (CloudNet), 2014 IEEE 3rd International Conference on* (pp. 138 - 142). IEEE. <https://doi.org/10.1109/CloudNet.2014.6968982>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

On the Throughput and Energy Benefits of Network Coded Cooperation

Néstor J. Hernández Marcano*, Janus Heide*, Daniel E. Lucani[†] and Frank H.P. Fitzek[†]

*Steinwurf ApS, Aalborg, Denmark. Mail: {nestor|janus}@steinwurf.com

[†]Department of Electronic Systems, Aalborg University, Denmark. Mail: {nh|del|ff}@es.aau.dk

Abstract—Cooperative techniques in wireless mobile networks typically leverage short-range communication technologies, e.g., WiFi, to allow data exchange between devices forming a mobile cloud. These mobile clouds have been considered as a key to reduce the cost of multicast services for the network operators as well as a means to deliver a better quality to the users. In fact, LTE-A includes Device-to-Device communication capabilities to enable such a direct communication between devices. The underlying assumption for attaining the throughput gains in mobile clouds is that the communication rate between devices is typically larger than the data rate from the base station to a receiver. However, while the data rates on cellular technologies have been steadily increasing, short-range communication speeds have remained largely unchanged calling into question these assumptions. This work’s goal is to assess the operating regions where the use of cooperation results in a higher throughput and/or energy saving. We consider a multicasting and a cooperative scheme with network coded mechanisms, as they typically outperform uncoded approaches. Our analysis and numerical results show that gains of several fold can be attained even if the data rate of the short-range technologies is moderately larger, e.g., 2x larger, than the cellular link data rate.

Keywords—4G, cooperation, energy, mobile clouds, network coding, throughput

I. INTRODUCTION

Data traffic is expected to grow by an order of magnitude for wireless mobile devices due largely to video services. This presents significant technical challenges for mobile operators to provide high quality of experience to the network users at high data rates, with low delay, while maintaining a low energy consumption in the mobile devices. Thus, mechanisms that can offload infrastructure networks have gathered significant interest from both academia and industry.

To address some of these challenges in multicast transmissions, wireless cooperation between receivers leveraging a separate communication channel to exchange missing data packets (instead of requesting them directly from the cellular infrastructure as in Fig. 1) are known to provide large gains over simply broadcasting the data [1].

This potential for cooperation has resulted in the inclusion of device-to-device (D2D) communication in the 3rd Generation Partnership Project (3GPP) standardization efforts. Beyond offloading the network operator, these cooperative techniques can result in increased reliability, coverage extension, and even increased throughput to end receivers.

In this context, network coding (NC) [2] provides not only a faster and more efficient approach to broadcast the data to the users, as shown by [3], but it simplifies the cooperation process

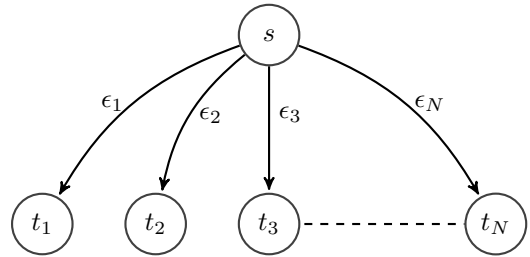


Fig. 1: Single source multiple sink topology of N receivers

since (i) devices need not know the specific packets missing at other devices, only the number of linear combinations available; and (ii) transmissions from a single device during the cooperation process can be used to *heal* a packet at multiple receivers, i.e., each transmission in the cooperation phase can have a larger impact for the end-receivers. This intuition has been exploited in previous work ranging from analysis to optimal policies and practical mechanisms, e.g., [4], [5].

However, the conventional wisdom of such cooperative techniques is that the secondary channel is considerably faster than the channel to the base station. Although this assumption was reasonable in the context of 2G and 3G communications, the much higher data rates achievable in LTE-A (4G) calls this assumption into question. The reason is that alternative technologies for device to device communications, e.g., WiFi, may no longer be faster than LTE-A as their data rates have stayed moderately constant over time. Additionally, if the devices cooperate using D2D of LTE-A the data rate for cooperation will also be limited by the common channel and could be the same data rate in some cases. Thus, the goal of this paper is to revisit the problem of device cooperation focusing on the specific regions of operation where it can bring gains in throughput and energy.

Some of the analysis of mean performance for cooperative schemes has been carried out before, e.g., [4], [6], however this paper provides an in-depth study of the distributions of the number of transmitted packets of different broadcast and cooperative schemes with NC (Section II). Leveraging these distributions, we derive the throughput and energy performance of the various schemes (Section III). In particular, we introduce the natural concept of stable throughput for cooperative schemes. To the best of our knowledge, this has not been considered before because of the conventional assumption that the cellular data rate is the bottleneck in the communication process. Our analysis allows us to determine the regions where cooperation provides gains over broadcasting (Section IV).

II. MODEL AND TRANSMISSION SCHEMES

We consider the problem of reliably transmitting a batch of packets from a source to N receivers using various transmission schemes. The batch constitutes a generation of g packets which we code using random linear network coding (RLNC) with field size q [7]. We assume independent heterogeneous erasure rates on the links from the source to the receivers, $\epsilon_j, j = [1, N]$, to derive the expressions.

We review two transmission schemes namely broadcast RLNC and cooperation with NC. For the cooperative scheme the receivers communicate among themselves to locally repair missing packets. We model the number of transmission as random variables using the geometric distribution as a building block to derive the probability mass functions (pmf) in order to obtain a complete description of the transmission process. We first give a new expression of the pmf for RLNC with no erasures and then compute the pmf for the schemes.

A. RLNC Probability Mass Function

Consider the case of a single source - destination pair without erasures. Let $\mathbb{T}_{RLNC,i}$ be a r.v. for the number of transmissions needed to receive a linearly independent (l.i.) coded packet in a stage of RLNC, i.e. once $i - 1$ l.i. packets have been received. This is a geometric distribution with success probability given by $p_i = 1 - q^{-g+(i-1)}, i \in [1, g]$. Following, $\mathbb{T}_{RLNC} = \sum_{i=1}^g \mathbb{T}_{RLNC,i}$ transmissions are necessary to decode g packets. Therefore, the code pmf can be computed using a characteristic function approach to make the analysis more tractable. Consequently, we obtain the pmf for RLNC without erasures in (1), where $P_g = \prod_{i=1}^g p_i = \Pr[\mathbb{T}_{RLNC} = g]$ is the probability of decoding in exactly g transmissions, $\gamma_i = 1 - p_i$ is the probability of receiving a linearly dependent coded packet and $a_i = \prod_{m=1, m \neq i}^g (1 - q^{m-i})^{-1}$ is a scaling factor for γ_i that quantifies the effect of the linear dependence in the decoding probability.

$$f_{\mathbb{T}_{RLNC}}(t; q, g) = \Pr[\mathbb{T}_{RLNC} = t] = P_g \sum_{i=1}^g a_i \gamma_i^{t-g}, \quad t \in [g, \infty) \quad (1)$$

B. Broadcast RLNC

As an approximation, we consider the case of finding the required transmissions for the maximum of N independent unicast sessions which makes the results an upper bound since we are excluding the transmissions accounting common coded packets. For each unicast session, we first model a single source - destination pair with erasure ϵ with RLNC and then proceed to calculate the broadcast case. Here, we need to account for g l.i. received packets in t transmissions. Hence, we need to consider all the cases where i l.i. packets are received (with the final success in t , which [8], [9] do not consider) and $t - i$ packets were lost or linearly dependent. For this, we review two main probabilities in the same way as [10]. First, let $\Pr[\mathbb{T}_{S_i} = t]$ be the probability for receiving i coded packets in t transmission (only considering the erasures), then $\mathbb{T}_{S_i} \sim \text{NB}(i, 1 - \epsilon)$. Second, the probability that g coded packets are l.i. in i slots, is $\Pr[\mathbb{T}_{RLNC} = i]$. Subsequently,

the probability of decoding in exactly t slots for a single user with RLNC based unicast with erasure ϵ , $\mathbb{T}_{U,cod}$, is:

$$f_{\mathbb{T}_{U,cod}}(t; \epsilon, q, g) = \Pr[\mathbb{T}_{U,cod} = t] = \sum_{i=g}^t \binom{t-1}{i-1} (1 - \epsilon)^i \epsilon^{t-i} f_{\mathbb{T}_{RLNC}}(i; q, g), \quad t \in [g, \infty) \quad (2)$$

Since each receiver just needs to collect different linear combinations to decode the packets, the number of transmissions will be bounded by the receiver that performs the worst in terms of retransmissions, i.e. $\mathbb{T}_{B,cod} = \max_{j=[1,N]} \mathbb{T}_{U_j,cod}$, which we calculate by a c.d.f. approach. For the probability of the maximum being less than or equal to t transmissions, this must occur for every receiver. Then, under the independence assumption, we can compute the c.d.f. for broadcast RLNC, e.g. $\Pr[\mathbb{T}_{B,cod} \leq t] = \prod_{j=1}^N \Pr[\mathbb{T}_{U_j,cod} \leq t]$ from (2) with the resulting pmf in (3).

$$f_{\mathbb{T}_{B,cod}}(t; N, \epsilon_1, \dots, \epsilon_N, q, g) = \prod_{j=1}^N \left(\sum_{k=g}^t \sum_{i=g}^k \binom{k-1}{i-1} (1 - \epsilon_j)^i \epsilon_j^{k-i} f_{\mathbb{T}_{RLNC}}(i; q, g) \right) - \prod_{j=1}^N \left(\sum_{k=g}^{t-1} \sum_{i=g}^k \binom{k-1}{i-1} (1 - \epsilon_j)^i \epsilon_j^{k-i} f_{\mathbb{T}_{RLNC}}(i; q, g) \right), \quad t \in [g, \infty) \quad (3)$$

We notice that the expression in (3) is the general case for low field sizes of the randomized broadcast coding scheme reviewed in [3], since if we let $q \rightarrow \infty$ in (3), then the c.d.f. used to compute (3) tends to the c.d.f. used to compute the mean and variance in section III-B of [3].

C. Cloud Cooperation with Coding

For the cooperation scheme, we consider a mobile cloud composed of H receivers ($H < N$) with cellular connection, the *heads* and $N - H$ receivers the *non-heads*, with a local connection to the heads. Packet transmissions takes place in two stages: (i) between source and heads and (ii) between heads and non-heads, which we label the cellular and local stage respectively. For the cellular stage, the source broadcast a coded packet to the heads which receive it *collectively*, i.e. it is enough that one head gets it for the cloud to acknowledge reception, with the stage finishing once the heads get the generation as a group. In the local stage, the heads broadcast recoded packets in a round robin fashion to the non-heads. The local stage finishes once all receivers have decoded the generation.

Under this condition, the distribution of receiving g packets in the cellular stage for the heads, $\mathbb{T}_{C,cod}$ is modeled as $\mathbb{T}_{U,cod}$ but with a success probability given by $1 - \prod_{j=1}^H \epsilon_j$ because all links need to fail for a packet to not be received. In the local stage the heads takes turn to broadcast to the non-heads, which is a particular case of (3). The total number of transmissions for this scheme, $\mathbb{T}_{CC,cod}$, is given by $\mathbb{T}_{CC,cod} = \mathbb{T}_{C,cod}(H, \prod_{j=1}^H \epsilon_j, g, q) + \mathbb{T}_{B,cod}(N - H, \epsilon, g, q)$ where the parentheses notation indicates the evaluation of the pmf with the given parameters.

III. PERFORMANCE METRICS

With the pmf for each scheme from section II, we calculate the moments for the number of transmissions which allows us to compute the throughput and energy.

A. Throughput

We define the throughput in the cloud cooperation scheme for a given erasure rate, generation and field size in the following way:

$$R_{eff,CC} = \frac{g}{\max(T_{s,cel}\mathbb{T}_{C,cod}(H), T_{s,loc}\mathbb{T}_{B,cod}(N-H))} \quad (4)$$

In (4), $T_{s,cel}$ and $T_{s,loc}$ are the duration of a time slot in the cellular and local stages, respectively. The effective rate perceived by a user will be the information sent divided by the completion time. For broadcast RLNC, the throughput is $R_{eff,B} = \frac{g}{T_{s,cel}\mathbb{T}_{B,cod}(N)}$.

B. Energy Consumption

We review the energy spent for the BS and average energy per receiver for the cooperation and broadcast schemes on the coded cases for a given erasure and code parameters. First, the energy consumption for broadcast is as follows:

$$E_{T_x} = E_{cel}\mathbb{T}_{B,cod}(N) ; E_{R_x} = E_{cel}\mathbb{T}_{B,cod}(N) \quad (5)$$

Where $E_{cel} = N_B E_B$ is the energy cost per packet in the cellular stage, N_B is the number of bytes per packet and E_B is the energy per byte proportional to the energy per bit. In a similar way, the energy expenditure for the cooperation schemes is shown in (6).

$$\begin{aligned} E_{T_x} &= E_{cel}\mathbb{T}_{C,cod}(H) \\ E_{R_x} &= E_{cel}\left(\frac{H}{N}\right)\mathbb{T}_{C,cod}(H) + E_{loc}\mathbb{T}_{B,cod}(N-H) \end{aligned} \quad (6)$$

C. Cellular vs. Local Links

The performance of cooperation will depend on the throughput and energy use per bit on the local links vs. that on the cellular links. Therefore, we define the r_t as the ratio between cellular and local throughput, and r_e as the ratio between the cellular and local energy.

$$r_t = \frac{T_{s,loc}}{T_{s,cel}} = \frac{R_{s,cel}}{R_{s,loc}} ; r_e = \frac{E_{b,cel}}{E_{b,loc}} \quad (7)$$

D. Gain Regions

For the analysis with different erasure rates per stage, we define the throughput and energy gains of cloud cooperation against broadcast RLNC from (5) and (6) as shown in (8).

$$\begin{aligned} G_t &= \frac{E\{\mathbb{T}_{B,cod}(N, \epsilon_{cel})\}}{\max(r_t E\{\mathbb{T}_{C,cod}(H, \epsilon_{cel})\}, E\{\mathbb{T}_{B,cod}(N, \epsilon_{cel})\})} \\ G_e &= 1 - \frac{r_e \left(\frac{H}{N}\right) E\{\mathbb{T}_{C,cod}(H, \epsilon_{cel}) + \mathbb{T}_{B,cod}(N-H, \epsilon_{loc})\}}{r_e E\{\mathbb{T}_{B,cod}(N, \epsilon_{cel})\}} \end{aligned} \quad (8)$$

We define throughput gain as the ratio of the cloud cooperation and broadcast RLNC throughputs. The energy gain of cooperation over broadcast is defined as the saving in energy for the devices, since cooperation always save energy at the BS.

IV. NUMERICAL RESULTS

With the obtained expressions, we can evaluate broadcast and cooperation to study the impact on the throughput and energy at the receivers, as we vary the number of users, the ratio between the cellular and local costs, and the erasure rates on the cellular and local links. We use a set of parameters in the following ranges $1 \leq N \leq 50$, $g = \{64, 128\}$, $q = 2^8$ and $0 \leq \epsilon \leq 0.6$. The timeslot duration is set to $T_{s,cel} = 0.5$ ms to conform to the LTE-A E-UTRA [11] and its set of D2D specifications. For the energy, we extracted the energy per bit cost from the energy model in [12] and use a packet size N_B of 500 B.

Fig. 2 shows the throughput as defined in (4) for the different cooperation schemes and broadcast RLNC when the cellular and local data rate are identical. Generally as the number of users increase the sustainable throughput to each receiver decreases. The highest throughput is obtained when the majority are heads, as this reduces the work in the local phase. As the number of non-heads increases the throughput with cooperation tends to that of broadcast, because the transmissions on the local stage becomes the dominating cost.

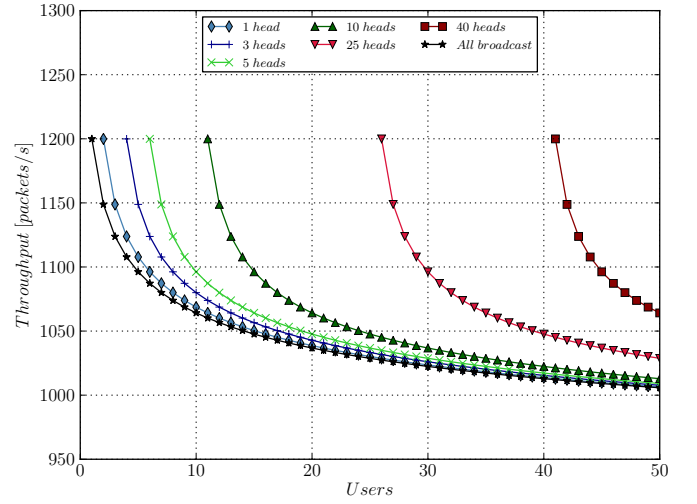


Fig. 2: Schemes throughput for equal data rate costs in the cellular and local link. Used parameters: $g = 64$, $q = 2^8$, $\epsilon = 0.4$, $r_t = 1$

Fig. 3 shows the energy spent per device where the energy costs are the same on the cellular and local links for both schemes. For a low amount of users, the energy consumption for the cooperation scheme is higher than broadcast because the amount of transmissions in the cellular and local links are comparable. As the number of user increases, the number of transmissions in the cellular link tends to g while the transmissions in the local link increases reducing the difference in performance.

Fig. 4 shows how the throughput varies depending on the ratio of the cellular and local data rate. The ratios are

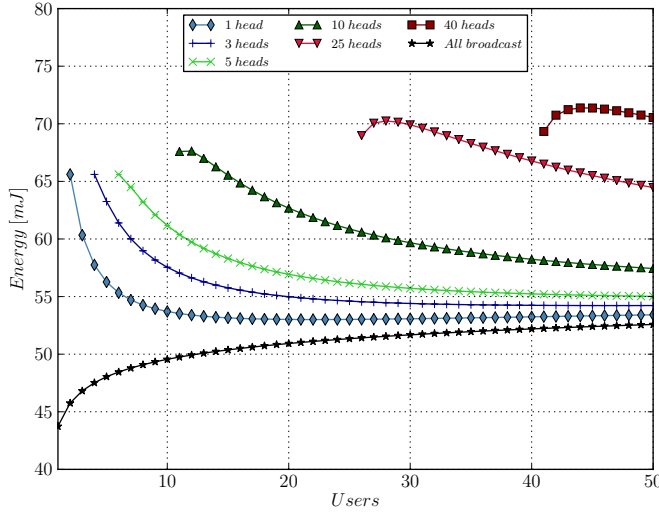


Fig. 3: Schemes energy consumption for equal energy costs in the cellular and local link. Used parameters: $g = 64$, $q = 2^8$, $\epsilon = 0.4$, $r_e = 1$

obtained by fixing the cellular data rate and varying the local data rate. When the local data rate is lower than the cellular rate the cooperative schemes provides lower throughput than the broadcast scheme. Conversely, when the local data rate is higher than the cellular data rate, the cooperative schemes delivers a higher throughput than broadcast. The throughput is highest when the local links rate are twice as faster as the cellular ones. The number of heads controls how much gain can be obtained and where it occurs for a given ratio. When the number of heads decreases, the throughput also diminishes because there are fewer heads each with an independent chance of receiving the packet.

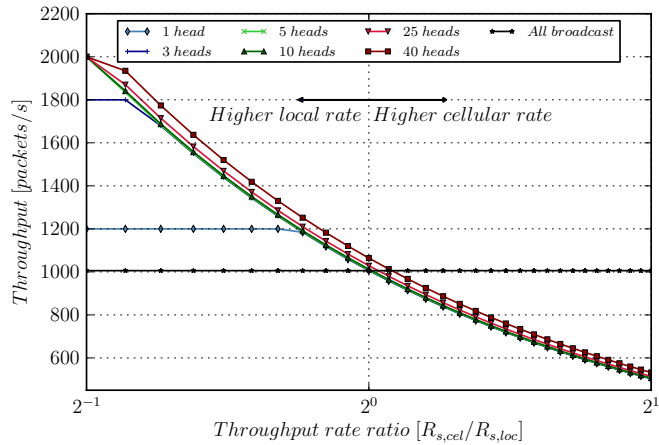


Fig. 4: The throughput of broadcast and cooperation with different number of heads, for different ratios between the data rate on the cellular and local link. Used parameters: $g = 64$, $q = 2^8$, $\epsilon = 0.4$, $N = 50$.

Fig. 5 shows how the energy for the devices changes as ratio between the cellular and local energy per bit changes. The energy cost in the cellular link is fixed and cost on the local link is changed to obtain the different ratios, consequently the energy per bit for broadcast is constant. When the energy

cost for the local links is higher than the cellular energy cost, the cooperative scheme expends more energy than the broadcast scheme. The additional consumption for cooperation comes from the transmissions in the local stage. Contrarily, when the cost of the local links is lower than the cost of the cellular links, then the cooperation scheme uses less energy than broadcast. For the cooperation schemes, the consumption is determined by the number of heads on the cellular stage. For a low number of heads, energy consumption is the lowest because the transmissions on the cellular links are for a few devices only.

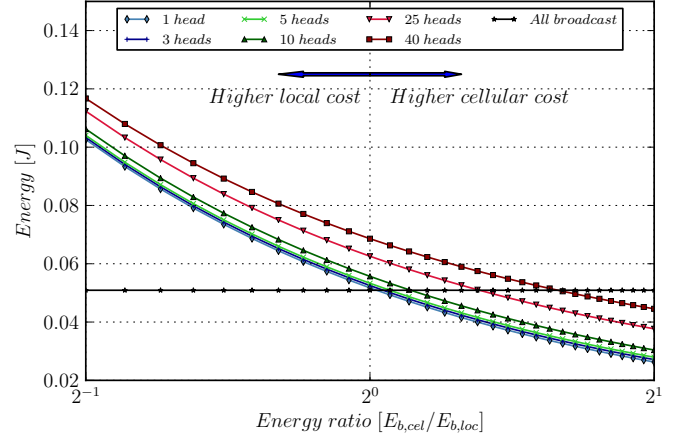


Fig. 5: The energy per generation of broadcast and cooperation with different number of heads, for different ratios between the energy per bit on the cellular and local link. Used parameters: $g = 64$, $q = 2^8$, $\epsilon = 0.4$, $N = 50$.

Fig. 6 shows the regions where cooperation provides a gain in terms of throughput for a wide range of erasure rates on the cellular and local links. The lines show where broadcast and cooperation performs the same, for $r_t = [0.5, 0.8, 1, 1.5, 2]$. In the region below each line, cooperation provides higher throughput than broadcast for that particular r_t . Above the line broadcast performs better. E.g. in the case of a fast local link $r_t = 0.5$ then cooperation provides a gain for almost all considered erasure rates, even in cases where the local erasure rate is much higher than the cellular.

Fig. 7 shows the regions where cooperation provides a gain in terms of energy saving on the devices for various erasure rates on the cellular and local links. The lines show where broadcast and cooperation performs the same, for $r_e = [0.5, 0.75, 1, 1.5, 2]$. In the region below each line, cooperation provides a lower energy per bit than broadcast for that particular r_e . Above the line broadcast performs better.

V. CONCLUSIONS

This work revisits the problem of wireless cooperation with network coding on cellular systems for multicast sessions in light of the increased data rates of current 4G and future 5G mobile networks and the stagnant data rates in short-range technologies, e.g., WiFi. This is particularly relevant because it breaks with the common assumption that the cooperative cluster can communicate locally at much higher data rates than the direct link to the cellular base station.

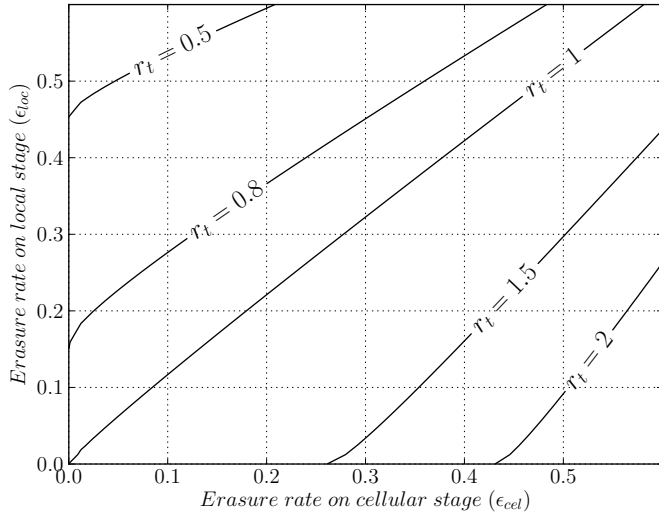


Fig. 6: For different values of r_t the lines indicate where cooperation and broadcast provide the same throughput for various erasure rates on the cellular and local links. Below the each line, cooperation performs better for the respective r_t . Used parameters: $g = 128$, $q = 2^8$, $H = 40$, $N = 50$.

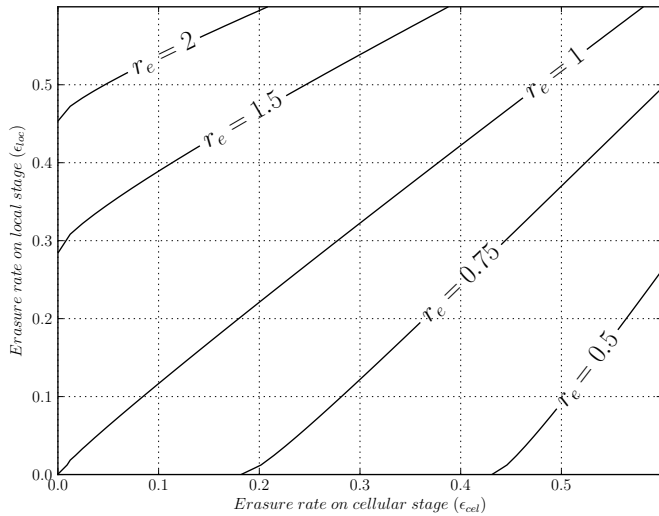


Fig. 7: For different values of r_e the lines indicate where cooperation and broadcast provide the same energy per bit for various erasure rates on the cellular and local links. Below the each line, cooperation performs better for the respective r_e . Used parameters: $g = 128$, $q = 2^8$, $H = 40$, $N = 50$.

More specifically, we presented an in-depth study of the specific operating regions where cooperation provides gains in throughput and energy over coded broadcasting techniques. Our numerical results showed that gains can be achieved even if the long-range and short-range technologies transmit at comparable data rates. More importantly, we showed that cooperation can provide several fold gains to the best broadcasting option (network coded broadcast) as long as the short-range link is at least twice as fast as the long-range one. Finally, our results showed that a moderate number of heads (e.g., three or more) per cooperative cluster is enough to yield the high throughput gains while maintaining a low energy consumption

at the receivers. The latter is not possible if a large fraction of the cooperative cluster is actively receiving directly from the base station with only a few exchanges needed during the cooperation process.

Future work shall focus on protocol design for cooperative schemes in highly-dense scenarios as well as implementation and evaluation of the most promising schemes in Aalborg University's Raspberry Pi testbed [13].

ACKNOWLEDGMENTS

This research has been partially financed by the CROSS-FIRE MITN Marie Curie project (317126) from the European Commission FP7 framework and the Green Mobile Cloud project (Grant No. DFF - 0602- 01372B) granted by the Danish Council for Independent Research.

REFERENCES

- [1] F. H. Fitzek and M. D. Katz, *Cooperation in Wireless Networks: Principles and Applications - Real Egoistic Behaviour is to Cooperate!* Springer, 2006.
- [2] R. Ahlswede, N. Cai, S.-Y. Li, and R. W. Yeung, "Network information flow," *Information Theory, IEEE Transactions on*, vol. 46, no. 4, pp. 1204–1216, 2000.
- [3] A. Eryilmaz, A. Ozdaglar, M. Médard, and E. Ahmed, "On the delay and throughput gains of coding in unreliable networks," *Information Theory, IEEE Transactions on*, vol. 54, no. 12, pp. 5511–5524, 2008.
- [4] J. Heide, F. H. Fitzek, M. V. Pedersen, and M. Katz, "Green mobile clouds: Network coding and user cooperation for improved energy efficiency," in *Cloud Networking (CLOUDNET), 2012 IEEE 1st International Conference on*. IEEE, 2012, pp. 111–118.
- [5] H. Khamfroush, D. E. Lucani, and J. Barros, "Minimizing the completion time of a wireless cooperative network using network coding," in *Personal Indoor and Mobile Radio Communications (PIMRC), 2013 IEEE 24th International Symposium on*. IEEE, 2013, pp. 2016–2020.
- [6] J. Heide, M. V. Pedersen, F. H. Fitzek, and T. Larsen, "Network coding for mobile devices-systematic binary random rateless codes," in *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*. IEEE, 2009, pp. 1–6.
- [7] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong, "A random linear network coding approach to multicast," *Information Theory, IEEE Transactions on*, vol. 52, no. 10, pp. 4413–4430, 2006.
- [8] O. Trullols-Cruces, J. M. Barcelo-Ordinas, and M. Fiore, "Exact decoding probability under random linear network coding," *Communications Letters, IEEE*, vol. 15, no. 1, pp. 67–69, 2011.
- [9] X. Zhao, "Notes on "exact decoding probability under random linear network coding"," *Communications Letters, IEEE*, vol. 16, no. 5, pp. 720–721, 2012.
- [10] D. E. Lucani, M. Médard, and M. Stojanovic, "Random linear network coding for time division duplexing: Field size considerations," in *Global Telecommunications Conference, 2009. GLOBECOM 2009. IEEE*. IEEE, 2009, pp. 1–6.
- [11] 36.913, "Requirements for further advancements for e-utra (lte-advanced)," 3GPP, Tech. Rep. Release 8, March 2009.
- [12] M. Lauridsen, L. Noël, T. B. Sørensen, and P. Mogensen, "An empirical lte smartphone power model with a view to energy efficiency evolution," *Intel Technology Journal*, vol. 18, no. 1, pp. 172–193, 2014.
- [13] A. Paramanathan, P. Pahlavan, S. Thorsteinsson, M. Hundebøll, D. Lucani, and F. Fitzek, "Sharing the pi: Testbed description and performance evaluation of network coding on the raspberry pi," in *2014 IEEE 79th Vehicular Technology Conference*, 2014.