

Enhancing Fog Radio Access Networks via Coded Multicast Fronthauling and Edge Caching for Multi-Connectivity Transmission

Hyunwoo Kim, Alessandro Fasano, Jaehwan Kim, and David Tse

*1*Department of Electrical and Computer Engineering, Pohang University of Science and Technology, Pohang-si, Gyeongbuk-do, 37673 Korea;

Abstract—This work studies the advantages of coded multicasting for the downlink of a Fog Radio Access Network (F-RAN) system equipped with a multicast fronthaul link. In this system, a control unit (CU) in the baseband processing unit (BBU) pool is connected to distributed edge nodes (ENs) through a multicast fronthaul link of finite capacity, and the ENs have baseband processing and caching capabilities. Each user equipment (UE) requests a file in a content library which is available at the CU, and the requested files are served by the closest ENs based on the cached contents and on the information received on the multicast fronthaul link. The performance of coded multicast fronthauling is investigated in terms of the delivery latency of the requested contents under the assumption of pipelined transmission on the fronthaul and edge links and of single-user encoding and decoding strategies based on the hard transfer of files on the fronthaul links. Extensive numerical results are provided to validate the advantages of the coded multicasting scheme compared to uncoded unicast and multicast strategies.

I. INTRODUCTION

Fog Radio Access Network (F-RAN) is a wireless cellular system that enables content delivery to user equipments (UEs) by means of both edge caching and cloud processing [1][3]. In an F-RAN, edge nodes (ENs), such as small-cell base stations (SBSs), can pre-fetch frequently requested, or popular, contents for storage in their local caches, while retrieving uncached information from the cloud. Prior works [4]-[6] studied the design of F-RANs from an informationtheoretic viewpoint. Instead, the references [7]-[9] took a signal processing perspectives by focusing on the design of beamforming strategies under different criteria, such as the delivery latency [7], the network cost [8] and the delivery rate [9].

All these papers assumed that every EN has a dedicated orthogonal fronthaul link to a control unit (CU) in the cloud.

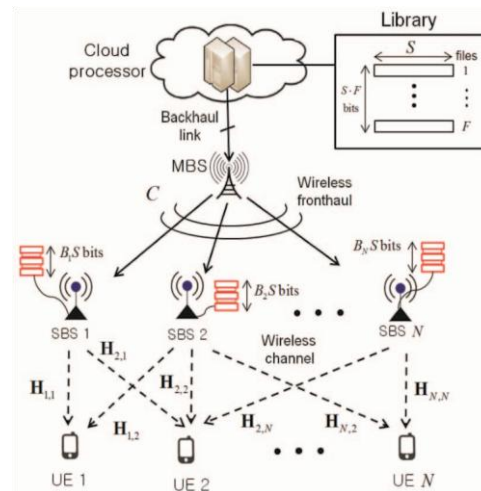


Figure 1. Illustration of an F-RAN model with wireless fronthaul link in which a cloud processor is connected to small-cell BSs, each equipped with local cache, through a macro BS which communicates with the small-cell BSs on a wireless fronthaul link.

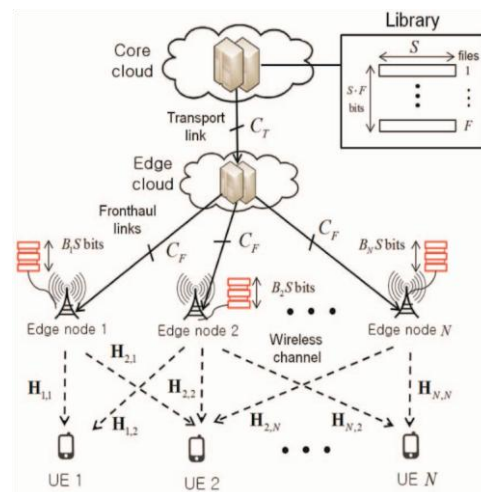


Figure 2. Illustration of a hierarchical F-RAN model, in which a core cloud has access to the content library and is connected to edge nodes, each equipped with local cache, through an edge cloud unit.

In contrast, in this work, we study an F-RAN system characterized by a shared multicast fronthaul link from the CU to the ENs. This model is motivated by two important deployment scenarios, which are shown in Fig. 1 and Fig. 2. In Fig. 1, the cloud processor is connected via a high-speed backhaul link to a macro BS (MBS) that provides a wireless fronthaul connection to a number of SBSs [10]. In the scenario of Fig. 2, the cloud has a



hierarchical structure with an *edge cloud* (EC) being connected to a *core cloud* (CC). Since the CC is connected to the ENs only through the EC, assuming that the latter only forwards the signal received from the CC, the link from the CC to the ENs can be equivalently modeled by a multicast link of capacity $C = \min\{C_T, C_F\}$, where C_T and C_F are the capacities of the CC-to-EC transport link and of the EC-to-ENs fronthaul links, respectively.

This work studies the latency needed for the delivery of the files requested by UEs for an F-RAN with a shared multicast fronthaul link. In particular, we investigate the advantages of coded multicasting [11] for fronthaul transmission by focusing on single-user encoding and decoding strategies based on the hard transfer of files on the fronthaul links [9]. With regards to communication from the ENs to the UEs on the edge wireless channels, we consider general multi-connectivity strategies as determined by the cached contents and focus on the design of linear beamforming strategies. We derive an efficient optimization solution based on the concave convex procedure (CCCP) and provide extensive numerical results that validate the advantages of coded multicasting as compared to conventional uncoded unicasting and multicasting approaches. *Notation:* We denote the mutual information between the random variables X and Y as $I(X;Y)$, and the circularly symmetric complex Gaussian distribution with mean μ and covariance matrix \mathbf{R} as $\mathcal{CN}(\mu, \mathbf{R})$. $\mathbb{C}^{M \times N}$ denotes the set of all $M \times N$ complex matrices and $\mathbb{E}(\cdot)$ represents the expectation operator. The operations $(\cdot)^T$ and $(\cdot)^\dagger$ denote the transpose and Hermitian transpose of a matrix, and $\lfloor x \rfloor$ is the largest integer not larger than x . The determinant and trace of a matrix \mathbf{X} are denoted as $\det(\mathbf{X})$ and $\text{tr}(\mathbf{X})$, respectively.

II. SYSTEM MODEL

As seen in Fig. 1, we consider the downlink of an F-RAN with N pairs of ENs and UEs, where the ENs are connected to a CU by means of a shared multicast fronthaul link of capacity C bit/symbol. Henceforth, a symbol refers to a channel use of the downlink wireless channel. In this system, the UEs request files from a static library of F popular files, where each file $f \in \mathcal{F} \triangleq \{1, \dots, F\}$ is of size S bits. It is assumed that the CU has access to the library and that the probability $p(f)$ of a file f to be selected is given by Zipf's distribution $p(f) = cf^{-\gamma}$ for $f \in \mathcal{F}$, where $\gamma \geq 0$ is a given popularity exponent and $c \geq 0$ is set such that $\sum_{f \in \mathcal{F}} p(f) = 1$. The file requested by the k th UE is denoted by $f_k \in \mathcal{F}$, which is assumed to be independent across the index k . We define

the demand vector $\mathbf{f} = [f_1 f_2 \dots f_N]^T$ and the set $\mathcal{F}_{\text{req}} = \cup_{k \in \mathcal{N}} \{f_k\}$ of requested files.

Each EN can cache $B_i S$ bits from the library, and we define the fractional caching capacity μ_i of EN i as

$$\mu_i = \frac{B_i}{F}. \tag{1}$$

We focus on the symmetric case of $B_i = B$ and $\mu_i = \mu$ for all $i \in \mathcal{N} \triangleq \{1, \dots, N\}$, but the discussion can be extended to general cases.

A. Channel model

We assume that each EN and UE are equipped with n_T and n_R antennas, respectively. Under a flat-fading channel model, the baseband signal $\mathbf{y}_k \in \mathbb{C}^{n_R \times 1}$ received by UE k in each transmission interval is given as

$$\mathbf{y}_k = \sum_{i \in \mathcal{N}} \mathbf{H}_{k,i} \mathbf{x}_i + \mathbf{z}_k = \mathbf{H}_k \mathbf{x} + \mathbf{z}_k, \tag{2}$$

where $\mathbf{x}_i \in \mathbb{C}^{n_T \times 1}$ is the baseband signal transmitted by EN i ; $\mathbf{H}_{k,i} \in \mathbb{C}^{n_R \times n_T}$ denotes the channel response matrix from EN i to UE k ; $\mathbf{z}_k \in \mathbb{C}^{n_R \times 1}$ is the additive noise distributed as $\mathbf{z}_k \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$; $\mathbf{H}_k \triangleq [\mathbf{H}_{k,1} \dots \mathbf{H}_{k,N}] \in \mathbb{C}^{n_R \times N n_T}$ collects the channel matrices $\mathbf{H}_{k,i}$ from all the ENs to UE k ; and $\mathbf{x} [\mathbf{x}_1^\dagger \dots \mathbf{x}_N^\dagger]^\dagger \in \mathbb{C}^{N n_T \times 1}$ is the signal transmitted by all the ENs. We assume that each EN i is subject to the average transmit power constraint $\mathbb{E} \|\mathbf{x}_i\|^2 \leq P$, and the signal-to-noise ratio (SNR) of the wireless edge link is defined as P . Furthermore, the channel matrices $\{\mathbf{H}_{k,i}\}_{k,i \in \mathcal{N}}$ are assumed to remain constant during each transmission interval. In Sec. V, we will evaluate the system performance under the additional assumption that the elements of the channel matrix $\mathbf{H}_{k,i}$ are independent and identically distributed (i.i.d.) as $\mathcal{CN}(0, \alpha^{k-i})$, where the parameter $0 < \alpha < 1$ accounts for the path loss. Note that this choice reflects the facts that UE k is closest to its serving EN k and that the distance between UE k and EN i increases with the difference $|k - i|$.

B. System Operation and Delivery Latency

The F-RAN system under study operates in two phases, namely pre-fetching and delivery, as in [9] and references therein (see [6] for an online problem formulation). In the prefetching phase, which takes place offline, say at night, each EN i populates its cache from the library of F files. In the following delivery phase, which spans many time slots, for any given demand vector \mathbf{f} , the CU and ENs cooperate to serve the UEs via the multicast fronthaul and wireless edge links in each time slot.

For the pre-fetching phase, we consider the randomized fractional cache distinct strategy studied in [9, Sec. III-C], in which each EN populates its cache in a distributed manner. With this strategy, each file f is split into L equal-

sized subfiles $(f,1), \dots, (f,L)$, such that each subfile (f,l) is of size $S^* = S/L$ bits. Each EN i stores randomly chosen $\tilde{L} = \lfloor \mu L \rfloor$ fragments of every file $f \in \mathcal{F}$. We define binary caching variables $\mathbf{c} = \{c_{f,l}^i\}_{f \in \mathcal{F}, l \in \mathcal{L}}$, with $\mathcal{L} = \{1, \dots, L\}$, as

$$c^i = \begin{cases} 1, & \text{if subfile } (f,l) \text{ is cached by EN } i \\ 0, & \text{otherwise} \end{cases} \quad f,l \quad (3)$$

which must satisfy the cache memory constraint

$$\sum_{f \in \mathcal{F}} \sum_{l \in \mathcal{L}} c_{f,l}^i \tilde{S} \leq BS, \quad (4)$$

for each EN i .

In the delivery phase, the CU and the ENs cooperate to deliver the requested files \mathcal{F}_{req} to the UEs. As in [7], we aim at designing the delivery strategies for the fronthaul and wireless edge links with the goal of minimizing the delivery coding latency. Specifically, we focus on single-user encoding and decoding strategies based on the hard transfer of files on the fronthaul links. This excludes interference management techniques such as dirty paper coding and cloudbased precoding as in cloud radio access network systems (see, e.g., [9, Sec. IV]). We also consider pipelined transmission on the fronthaul and edge links [4, Sec. VII], such that the overall delivery coding latency T_{total} is given as

$$T_{\text{total}} = \max\{T_F, T_E\}, \quad (5)$$

where T_F and T_E represent the coding latency required to communicate on the fronthaul and edge links, respectively. In the following sections, we describe the latency metrics T_E and T_F under multi-connectivity transmission and various multicasting strategies.

III. MULTI-CONNECTIVITY WIRELESS TRANSMISSION

For the efficient management of inter-UE interference signals on the wireless channel, we consider multi-connectivity transmission across the ENs such that UE k is served by the closest $M \leq N$ ENs, i.e.,

$$\mathcal{N}_{\text{EN},k} = \left\{ k - \left\lfloor \frac{M-1}{2} \right\rfloor, k - \left\lfloor \frac{M-1}{2} \right\rfloor + 1, \dots, k + \left\lfloor \frac{M}{2} \right\rfloor \right\}, \quad (6)$$

where M is referred to as connectivity level (see, e.g., [12]). Note that, in order for this approach to be implemented, the ENs in $\mathcal{N}_{\text{EN},k}$ should have the information of the content f_k requested by UE k either by means of caching or via the information received on the multicast fronthaul link.

Accordingly, the vector \mathbf{x} transmitted by the ENs can be written as

$$\mathbf{x} = \sum_{k \in \mathcal{N}} \mathbf{V}_k \mathbf{s}_k, \quad (7)$$

where $\mathbf{s}_k \in \mathbb{C}^{n_s \times 1}$ is the baseband signal vector that encodes the file f_k and is distributed as $\mathbf{s}_k \sim \text{CN}(\mathbf{0}, \mathbf{I})$; and $\mathbf{V}_k = [\mathbf{V}_{k,1}^\dagger \dots \mathbf{V}_{k,N}^\dagger]^\dagger \in \mathbb{C}^{N n_T \times n_s}$ is the precoding matrix for the

signal \mathbf{s}_k with the submatrix $\mathbf{V}_{k,i} \in \mathbb{C}^{n_T \times n_s}$ corresponding to EN i . Note that $n_s \leq \min\{N n_T, n_R\}$ represents the number of data streams that encode any file f_k , and that the precoding matrices $\mathbf{V}_{k,i}$ are subject to the connectivity condition

$$\text{tr}(\mathbf{V}_{k,i} \mathbf{V}_{k,i}^\dagger) = 0, \text{ for all } (k,i) \text{ with } i \notin \mathcal{N}_{\text{EN},k}. \quad (8)$$

This equality imposes that the file \mathbf{s}_k cannot be precoded by EN i unless the EN belongs to the set $\mathcal{N}_{\text{EN},k}$.

Under the precoding model (7) and the assumption that each UE k decodes the file f_k based on the received signal \mathbf{y}_k in

(2) by treating interference as noise, the achievable rate R_k for UE k is given as

$$R_k = g_k(\mathbf{V}) \triangleq I(\mathbf{s}_k; \mathbf{y}_k) \quad (9)$$

$$= \phi \left(\mathbf{H}_k \mathbf{V}_k \mathbf{V}_k^\dagger \mathbf{H}_k^\dagger, \sum_{l \in \mathcal{N} \setminus \{k\}} \mathbf{H}_k \mathbf{V}_l \mathbf{V}_l^\dagger \mathbf{H}_k^\dagger + \mathbf{I} \right),$$

where we defined the notation $\mathbf{V} = \{\mathbf{V}_k\}_{k \in \mathcal{N}}$ and the function $\phi(\mathbf{A}, \mathbf{B}) \triangleq \log_2 \det(\mathbf{A} + \mathbf{B}) - \log_2 \det(\mathbf{B})$. For given delivery rates $\mathbf{R} = \{R_k\}_{k \in \mathcal{N}}$, the latency T_E on the edge link is given as

$$T_E = \frac{S}{\min_{k \in \mathcal{N}} R_k}, \quad (10)$$

since $\min_{k \in \mathcal{N}} R_k$ is the rate at which all requested files \mathcal{F}_{req} can be delivered to the UEs. We note that increasing the connectivity level M always improves the rates R_k and thus reduces the edge latency T_E by relaxing the constraints (8). However, we will see in Sec. IV that this does not guarantee improved total latency T_{total} due to the increased fronthaul overhead.

From (10), we can see that the problem of minimizing the latency on the edge link is equivalent to that of maximizing the minimum rate $R_{\text{min}} \triangleq \min_{k \in \mathcal{N}} R_k$. Thus, we consider the problem:

$$\text{maximize } R_{\text{min}} \quad (11a)$$

$$\text{s.t. } R_{\text{min}} \leq g_k(\mathbf{V}), k \in \mathcal{N}, \quad (11b)$$

$$\text{tr}(\mathbf{V}_{k,i} \mathbf{V}_{k,i}^\dagger) \leq 0, k \in \mathcal{N}, i \notin \mathcal{N}_{\text{EN},k}, \quad (11c)$$

$$\sum_{k \in \mathcal{N}} \text{tr}(\mathbf{V}_{k,i} \mathbf{V}_{k,i}^\dagger) \leq P, i \in \mathcal{N} \quad (11d)$$

To tackle the non-convex problem (11), as in [9], we restate the problem with respect to the variables $\tilde{\mathbf{V}}_k = \mathbf{V}_k \mathbf{V}_k^\dagger$ and

relax the constraint $\text{rank}(\mathbf{V}^{\sim}_k) \leq n_s$. Then, we obtain a difference-of-convex problem and hence can derive an iterative algorithm based on the CCCP approach that gives non-decreasing objective values with respect to the number of iterations (see, e.g., [8]). After convergence, we obtain the precoding matrices \mathbf{V}_k by taking the n_s leading eigenvectors of \mathbf{V}^{\sim}_k multiplied by the square roots of the corresponding eigenvalues. Details follow as in [9].

IV. FRONTHAUL DELIVERY STRATEGIES

In this subsection, we discuss fronthauling strategies on the multicast fronthaul link and derive the corresponding latency metrics. To elaborate, we define a binary variable $d_{f_k,l}^i$ as

$$d_{f_k,l}^i = \begin{cases} 1 & i \in \mathcal{N}_{\text{EN},k} \text{ and } c_{f_k,l}^i = 0 \\ 0 & \text{otherwise} \end{cases}, \quad (12)$$

for $k, i \in \mathcal{N}$ and $l \in \mathcal{L}$. By the definition (12), if $d_{f,l}^i = 1$, the CU needs to send EN i the subfile (f,l) so as to enable multicasting transmission, since this subfile is not present in the cache of EN i . We discuss some baseline uncoded fronthauling strategies and then present the coded multicasting approach.

A. Uncoded Unicasting

With the baseline uncoded unicasting, the CU uses the fronthaul link to send each EN i the subfiles (f,l) with $d_{f,l}^i = 1$. In this approach, the overlap between the sets of subfiles needed by different ENs according to (12) is not taken into account. Therefore, the number S_B of bits transferred on the multicast link is given as

$$S_B = \sum_{i \in \mathcal{N}} \sum_{f \in \mathcal{F}_{\text{req}}} \sum_{l \in \mathcal{L}} d_{f,l}^i \tilde{S}, \quad (13)$$

and the latency T_F on the fronthaul link becomes

$$T_F = \frac{S_B}{C}. \quad (14)$$

B. Uncoded Multicasting

Since the multicast link is shared among the ENs, the subfiles needed by multiple ENs need not be transferred separately to each EN. The uncoded multicasting strategy hence transfers any subfile requested by multiple ENs only once to minimize the number of fronthaul usages. The number S_B of bits multicast on the fronthaul link with this approach is hence given as

$$S_B = \sum_{f \in \mathcal{F}_{\text{req}}} \sum_{l \in \mathcal{L}} 1 \left(\sum_{i \in \mathcal{N}} d_{f,l}^i > 0 \right) \tilde{S}, \quad (15)$$

where $1(\cdot)$ denotes the indicator function, which returns 1 if the argument statement is true or 0 otherwise. Thus, the fronthaul latency T_F is given as (14) with S_B given as (15).

C. Coded Multicasting

Since the CU communicates with the ENs, each equipped with cached contents, via a shared multicast link, coded multicasting can potentially reduce the fronthaul overhead. Accordingly, the CU sends coded subfiles obtained as linear combinations of the uncached subfiles in such a way that the intended ENs can decode the coded subfiles based on the cached subfiles (see, e.g., [11]). The number S_B of bits transferred on the fronthaul link in this approach is given as $S_B = n_{\text{sub}} S^{\sim}$, where n_{sub} is the number of coded subfiles that are transferred on the fronthaul link. This can be efficiently obtained by using the (suboptimal) greedy constrained local coloring algorithm proposed in [11, Sec. IV-A]. Having computed S_B via the algorithm, whose details can be found in [11, Sec. IV-A], the fronthaul latency T_F is given as (14).

As a final remark, we note that, for all the discussed fronthauling strategies, the fronthaul latency T_F increases with the connectivity level M due to the larger number of subfiles that need to be transferred on the fronthaul link. This suggests that the optimal connectivity level M should be carefully selected by considering its conflicting impacts on the edge and fronthaul latencies.

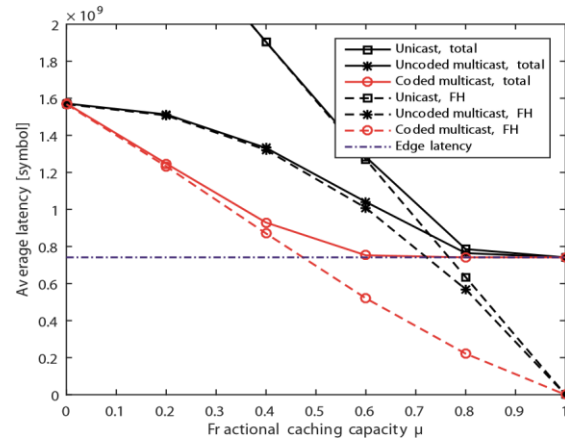


Figure 3. Average latency T_{total} versus the fractional caching capacity μ for an FRAN downlink ($F = 60, L = 50, N = 4, n_T = n_R = 1, C = 2, P = 20$ dB and $M = 2$).

V. NUMERICAL RESULTS

In this section, we present some numerical results to compare the latency of various fronthauling strategies discussed for the downlink of an F-RAN system with a shared multicast fronthaul link. Throughout the section,

we set $\gamma = 0.2$, $S = 100\text{MB}$ and $\alpha = 0.7$. We evaluate the total latency T_{total} , the fronthaul latency T_F and the edge latency T_E by averaging over the realizations of the caching variables \mathbf{c} , the UEs' requests \mathbf{f} and the channel matrices $\{\mathbf{H}_{k,i}\}_{k,i \in N}$.

In Fig. 3, we first investigate the impact of the fractional caching capacity μ on the average latency T_{total} , the average fronthaul latency T_F and the average edge latency T_E for an F-RAN downlink with $F = 60$, $L = 50$, $N = 4$, $n_T = n_R = 1$, $\alpha = 0.7$, $C = 2$, $P = 20$ dB and $M = 2$. Note that the edge latency T_E is the same for unicast and multicast fronthaul transmissions. It is observed that multicasting outperforms uncoded transmission, particularly for small values of μ , in which regime fronthaul latency determines the overall latency. Furthermore, coded multicasting can improve over uncoded multicasting, except for $\mu = 0$, in which case no side information is available at the ENs via caching, and for $\mu = 1$, in which case fronthaul transmission is unnecessary. The gain is seen to be due to a reduction in the fronthaul latency.

In Fig. 4, we plot the average latency T_{total} versus the connectivity level M for an F-RAN downlink with $F = 50$, $L = 20$, $N = 4$, $n_T = n_R = 1$, $P = 20$ dB and $\mu = 0.3$. The figure confirms that there is an optimum connectivity level M that strikes the best trade-off between fronthaul and edge latencies. It is also seen that the optimal value M increases with the fronthaul capacity C . Furthermore, the gain of coded multicasting is more relevant for a larger connectivity level M due to the increased coding opportunities.

A related conclusion can be reached from Fig. 5, which plots the average latency T_{total} versus the number L of fragments for an F-RAN downlink with $F = 60$, $N = 4$, $n_T = n_R = 2$, $P = 20$ dB, $C = 0.5$ and $M = 1$. The figure shows that

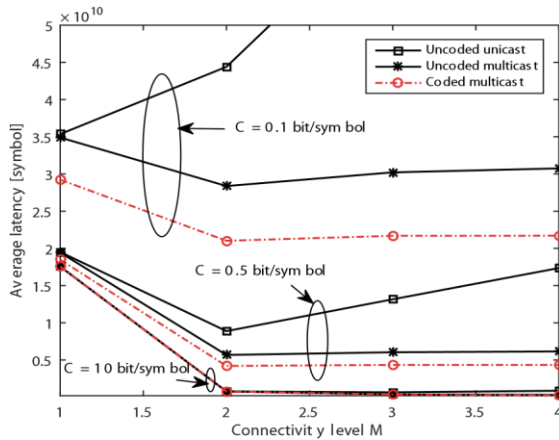


Figure 4. Average latency T_{total} versus the connectivity level M for an F-RAN downlink ($F = 50$, $L = 20$, $N = 4$, $n_T = n_R = 1$, $P = 20$ dB and $\mu = 0.3$).

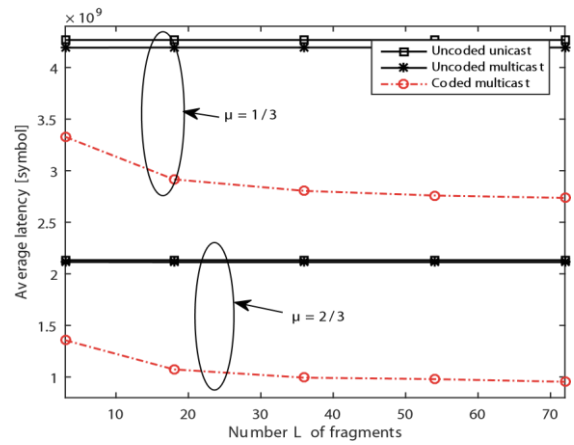


Figure 5. Average latency T_{total} versus the number L of fragments for an F-RAN downlink ($F = 60$, $N = 4$, $n_T = n_R = 2$, $P = 20$ dB, $C = 0.5$ and $M = 1$).

the gain of coded multicasting becomes more significant for a larger L due to the larger number of coding opportunities.

Lastly, in Fig. 6, we show the impact of the SNR P for an F-RAN downlink with $F = 60$, $L = 60$, $n_T = n_R = 1$, $\mu = 1/3$, $C = 1.0$, $N = 4$ and $M = 2$. The fronthaul latency does not change with the SNR on the edge channel, and hence the latency decrease is due to a reduction in the edge latency. As the SNR P increases, the total latency T_{total} becomes limited by the fronthaul latency while the edge latency dominates at lower SNR values. As a result, the gain of coded multicasting is observed in the regime of sufficiently large P .

VI. CONCLUSIONS

This paper studies the delivery coding latency for the downlink of an F-RAN system with a shared multicast fronthaul link. Under the assumption of pipelined transmission on the fronthaul and edge links, the advantages of coded multicast delivery on the fronthaul link were investigated for multiconnectivity transmission across the ENs and randomized fractional caching. We provided extensive numerical results

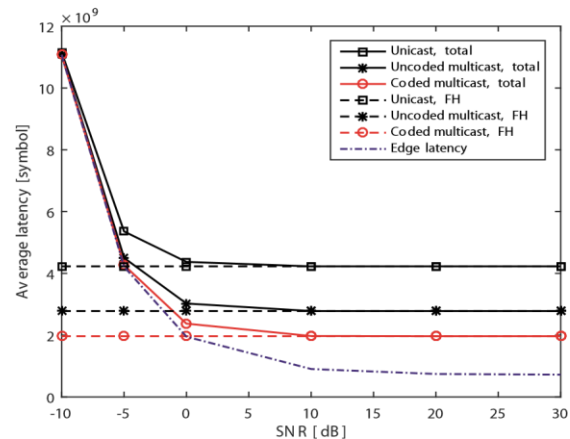


Figure 6. Average latency T_{total} versus the SNR P for an F-RAN downlink ($F = 60$, $L = 60$, $n_T = n_R = 1$, $\mu = 1/3$, $N = 4$, $C = 1.0$ and $M = 2$).

that validate the performance gains of the coded multicasting strategy as compared to the conventional uncoded strategies. Among open problems, we mention the development of an information-theoretic analysis that accounts for the potential performance gains of coded multicast fronthauling [10].

REFERENCES

- [1] S. Bi, R. Zhang, Z. Ding and S. Cui, "Wireless communications in the era of big data," *IEEE Comm. Mag.*, vol. 53, no. 10, pp. 190-199, Oct. 2015.
- [2] E. Zeydan, E. Bastug, M. Bennis, M. A. Kader, A. Karatepe, A. Salih Er and M. Debbah, "Big data caching for networking: Moving from cloud to edge," *IEEE Comm. Mag.*, vol. 54, no. 9, pp. 36-42, Sep. 2016.
- [3] M. Peng, Y. Sun, X. Li, Z. Mao and C. Wang, "Recent advances in cloud radio access networks: System architectures, key techniques, and open issues," *IEEE Comm. Surv. Tutorials*, vol. 18, no. 3, pp. 2282-2308, third quarter, 2016.
- [4] A. Sengupta, R. Tandon and O. Simeone, "Cloud and cache-aided wireless networks: Fundamental latency trade-offs," arXiv:1605.01690, May 2016.
- [5] R. Tandon and O. Simeone, "Cloud-aided wireless networks with edge caching: Fundamental latency trade-offs in fog radio access networks," in *Proc. IEEE Intern. Symp. Inf. Theory (ISIT) 2016*, Barcelona, Spain, Jul. 2016.
- [6] S. M. Azimi, O. Simeone, A. Sengupta and R. Tandon, "Online edge caching in fog-aided wireless network," arXiv:1701.06188, Jan. 2017.
- [7] S.-H. Park, O. Simeone and S. Shamaï (Shitz), "Joint cloud and edge processing for latency minimization in fog radio access networks," in *Proc. IEEE Intern. Workshop on Sig. Proc. Adv. Wireless Comm. (SPAWC) 2016*, Edinburgh, UK, Jul. 2016.
- [8] M. Tao, E. Chen, H. Zhou and W. Yu, "Content-centric sparse multicast beamforming for cache-enabled cloud RAN," *IEEE Trans. Wireless Comm.*, vol. 15, no. 9, pp. 6118-6131, Sep. 2016.
- [9] S.-H. Park, O. Simeone and S. Shamaï (Shitz), "Joint optimization of cloud and edge processing for fog radio access networks," *IEEE Trans. Wireless Comm.*, vol. 15, no. 11, pp. 7621-7632, Nov. 2016.
- [10] J. Koh, O. Simeone, R. Tandon and J. Kang, "Cloud-aided edge caching with wireless multicast fronthauling in fog radio access networks," in *Proc. IEEE Wireless Comm. Network. Conf. (WCNC) 2017*, San Francisco, CA, Mar. 2017.
- [11] M. Ji, A. M. Tulino, J. Llorca and G. Caire, "Order-optimal rate of caching and coded multicasting with random demands," arXiv:1502.03124, Feb. 2015.
- [12] A. Maeder, A. Ali, A. Bedekar, A. F. Cattoni, D. Chandramouli, S. Chandrasekar, L. Du, M. Hesse, C. Sartori and S. Turtinen, "A scalable flexible radio access network architecture for fifth generation mobile networks," *IEEE Comm. Mag.*, vol. 54, no. 11, pp. 16-23, Nov. 2016.