A New Interference-Alignment Scheme for Wireless MapReduce

Yue Bi[*][†], Michèle Wigger[†], Yue Wu[*]

[*]*School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University*, China; wuyue@sjtu.edu.cn
[†] *LTCI, Telecom Paris, IP Paris*, 91120 Palaiseau, France; {bi, michele.wigger}@telecom-paris.fr

Abstract—We consider a full-duplex wireless Distributed Computing (DC) system under the MapReduce framework. New upper and lower bounds on the optimal tradeoff between Normalized Delivery Time (NDT) and computation load are presented. The upper bound strictly improves over the previous reported upper bounds and is based on a novel interference alignment (IA) scheme tailored to the interference cancellation capabilities of MapReduce nodes. The lower bound is proved through information-theoretic converse arguments.

Index Terms—Wireless distributed computing, MapReduce, coded computing, interference alignment.

I. INTRODUCTION

Distributed Computing (DC) systems are computer networks that through task-parallelization reduce execution times of complex computing tasks such as data mining or computer vision. MapReduce is a popular such framework and runs in three phases [1], [2]. In the first *map phase*, nodes calculate intermediate values (IVA) from their associated input files. In the following *shuffle phase*, nodes exchange these IVAs in a way that each node obtains all IVAs required to compute its assigned output function in the final *reduce phase*. MapReduce is primarily applied to wired systems where it has been noticed that a significant part of the MapReduce execution time stems from the IVA *delivery time* during the shuffle phase [2], [3], and can be reduced through smart coding [3]–[8].

MapReduce systems are becoming important building blocks also in wireless scenarios, such as vehicular networks [9] or distributed e-health applications [10], thus creating a need for good wireless MapReduce coding schemes. Similarly to the wired case [4]–[6], delivery time in wireless MapReduce systems can be decreased by sending appropriate linear combinations of the IVAs, from which the receiving nodes can extract their desired IVAs by bootstrapping the IVAs that they can compute from their locally stored input files. Further improvements can however be achieved by exploiting specific wireless communication techniques.

The focus of this paper is on the high Signal-to-Noise Ratio (SNR) regime, and on the following two key metrics:

- *Computation load* r: This describes the average number of nodes to which each file is assigned. In other words, it is the ratio of the total number of assigned input files (including replications) normalized by the total number of files.
- *Normalized Delivery Time (NDT)* Δ: This is the wireless shuffle duration normalized by the number of reduce functions and input files and by the transmission time of a single IVA over a point-to-point channel.

We are interested in the minimal NDT for given computation load r, which we call the *NDT-computation tradeoff*.

The NDT of full-duplex interference networks was considered in [11], [12], see [13] for the half-duplex network. In [11], the authors proposed to divide the nodes into groups and let each group store a subset of the files and apply one-shot beamforming and zero-forcing during the shuffle communication. As shown in [11], their scheme is optimal among this class of communication strategies. In our previous work [12], we reduced this NDT by introducing the IA technique to the shuffle communication in [11].

In this paper, we obtain yet a further NDT reduction by considering the map procedure in [4], where *each set of* r *nodes* stores a subset of the files, and by proposing a novel IA scheme that is tailored to this file assignment and the interference cancellation capabilities of MapReduce nodes so as to obtain an improved performance compared to standard IA schemes. Our scheme is related to the IA scheme in [14], which considers a similar file assignment, and to our previous work [12], with which it coincides when each node can only store a single file.

The upper bound on the NDT implied by our new IA-scheme improves over the previously proposed bounds in [11], [12] whenever the computation load r satisfies $1 < r < \lceil \frac{K-1}{2} \rceil$. Our results thus show that in this regime, beamforming and zero-forcing cannot achieve minimum NDT. On the contrary, through an information-theoretic lower bound on the minimum NDT, in this manuscript we show that for $r \geq \frac{K}{2}$ the zero-forcing and interference cancellation scheme in [11] is optimal also without any restriction on the utilized coding scheme.

Notations: We use standard notation, and also define $[n] \triangleq \{1, 2, \ldots, n\}$ and $[\mathcal{A}]^n$ as the collection of all the subsets of $\mathcal{A}$ with cardinality n.

II. WIRELESS MAPREDUCE FRAMEWORK

Consider a distributed computing (DC) system with a fixed number of K nodes labelled $1, \ldots, K$; a large number N of input files $W_1, \ldots, W_N$; and K output functions $\phi_1, \ldots, \phi_K$ mapping the input files to the desired computations.

A *MapReduce* System decomposes the output functions as:

$$\phi_q(W_1, \ldots, W_N) = v_q(a_{q,1}, \ldots, a_{q,N}), \qquad q \in [K], \quad (1)$$

where v_q is called *reduce function* and $a_{q,p}$ the *intermediate value (IVA)* calculated from file W_p through a *map function*

$$a_{q,p} = u_{q,p}(W_p), \qquad p \in [\mathsf{N}]. \tag{2}$$

IVAs are independent with A i.i.d. bits.

The MapReduce framework has 3 phases:

Map phase: A subset of all input files $\mathcal{M}_k \subseteq [\mathsf{N}]$ is assigned to each node $k \in [\mathsf{K}]$. Node k computes all IVAs $\{a_{q,p}: p \in \mathcal{M}_k, q \in [\mathsf{K}]\}$ associated with these input files. Notice that the set $\{\mathcal{M}_k\}_{k \in [\mathsf{K}]}$ is a design factor.

Shuffle phase: Computation of the k-th output function is assigned to the k-th node.

The K nodes in the system communicate over T uses of a wireless network in a full-duplex mode to exchange the . missing IVAs for the computations of their assigned output functions. So, node $k \in [\mathsf{K}]$ produces complex channel inputs of the form

$$\mathbf{X}_k \triangleq (X_k(1), \dots, X_k(\mathsf{T}))^T = f_k^{(\mathsf{T})}\left(\{a_{1,p}, \dots, a_{\mathsf{K},p}\}_{p \in \mathcal{M}_k}\right), \tag{3}$$

by means of an encoding function $f_k^{(\mathsf{T})}$ on appropriate domains and so that the inputs satisfy the block-power constraint

$$\frac{1}{\mathsf{T}} \sum_{t=1}^{\mathsf{T}} \mathbb{E}\left[|X_k(t)|^2\right] \leq \mathsf{P}, \qquad k \in [\mathsf{K}]. \tag{4}$$

Given the full-duplex nature of the network, Node k also observes the complex channel outputs

$$Y_k(t) = \sum_{k' \in [\mathsf{K}]\setminus\{k\}} H_{k,k'}(t) X_{k'}(t) + Z_k(t), \quad t \in [\mathsf{T}], \tag{5}$$

where the sequences of complex-valued channel coefficients $\{H_{k,k'}(t)\}$ and standard circularly symmetric Gaussian noises $\{Z_k(t)\}$ are both i.i.d. and independent of each other and of all other channel coefficients and noises.

Based on its outputs $\mathbf{Y}_k \triangleq (Y_k(1), \dots, Y_k(\mathsf{T}))^T$ and the IVAs $\{a_{q,p}: p \in \mathcal{M}_k, q \in [\mathsf{K}]\}$ it computed during the Map phase, Node k decodes the missing IVAs $\{a_{k,p}: p \notin \mathcal{M}_k\}$ required to compute its assigned output functions ϕ_k as:

$$\hat{a}_{k,p} = g_{k,p}^{(\mathsf{T})}\left(\{a_{1,i}, \dots, a_{\mathsf{K},i}\}_{i \in \mathcal{M}_k}, \mathbf{Y}_k\right), \quad p \notin \mathcal{M}_k. \tag{6}$$

Reduce phase: Each node $k \in [\mathsf{K}]$ applies reduce functions $\phi_k(\cdot)$ to the appropriate IVAs calculated during the Map phase or decoded in the Shuffle phase.

The performance of the distributed computing system is measured in terms of its *computation load*

$$\mathsf{r} \triangleq \sum_{k \in [\mathsf{K}]} \frac{|\mathcal{M}_k|}{\mathsf{N}}, \tag{7}$$

and the *normalized delivery time (NDT)*

$$\Delta \triangleq \varlimsup_{\mathsf{P} \to \infty} \varlimsup_{\mathsf{A} \to \infty} \frac{\mathsf{T}}{\mathsf{A} \cdot \mathsf{K} \cdot \mathsf{N}} \cdot \log \mathsf{P}. \tag{8}$$

We focus on the *fundamental NDT-computation tradeoff* $\Delta^*(\mathsf{r})$, which is defined as the infimum over all values of Δ satisfying (8) for some choice of file assignments $\{\mathcal{M}_k\}$ and

sequence (in T) of encoding and decoding functions $\{f_k^{(\mathsf{T})}\}$ and $\{g_{k,p}^{(\mathsf{T})}\}$ in (3) and (6), all depending on A so that the probability of IVA decoding error

$$\Pr\left[\bigcup_{k \in [\mathsf{K}]} \bigcup_{p \notin \mathcal{M}_k} \hat{a}_{k,p} \neq a_{k,p}\right] \to 0 \quad \text{as} \quad \mathsf{A} \to \infty. \tag{9}$$

A. Sufficiency of Symmetric File Assignments

Our model exhibits a perfect symmetry between the various nodes in the network because the channels from any Tx-node to any Rx-node are independent and have identical statistics. The optimal NDT-computation tradeoff can therefore be achieved by a symmetric file assignment where any subset of nodes $\mathcal{T} \subseteq [\mathsf{K}]$ of size i is assigned the same number of files to be stored at all nodes in $\mathcal{T}$. In fact, any non-symmetric file assignment can be symmetrized without decreasing the NDT-computation tradeoff. It suffices to time-share K! instances of the original scheme for a IVA size that is also multiplied by K!, where in each instance the K nodes are relabeled according to a different permutation. The resulting scheme has a symmetric file assignment and achieves the same NDT-computation tradeoff as the original scheme because both T and A are multiplied by K! while the other parameters remain unchanged and because the new scheme still satisfies (9).

B. Relation to the Sum-DoF with r-fold Cooperation

A well-studied property of wireless networks is the *Sum Degrees of Freedom (sum-DoF)*, which characterizes the maximum throughput of a network. In this work we are specifically interested in the sum-DoF that one can achieve over the wireless network described by (5), when the inputs are subject to the average power constraints (4) and any set of r nodes $\mathcal{T} \in [\mathsf{K}]^r$ has a message $M_{\mathcal{T}}^j$ that it wishes to convey to Node j, for any $j \in [\mathsf{K}]\setminus\mathcal{T}$. Each message $M_{\mathcal{T}}^j$ is uniformly distributed over a set $[2^{nR_{\mathcal{T}}^j}]$ and a rate-tuple $(R_{\mathcal{T}}^j: \mathcal{T} \in [\mathsf{K}]^r, j \in [\mathsf{K}]\setminus\mathcal{T})$ is called achievable if there exists a sequence of encoding and decoding functions such that the probabilities of error tend to 0 in the asymptotic regime of infinite blocklengths. The sum-DoF is then defined as

$$\text{Sum-DoF}(\mathsf{r}) \triangleq \sup \varlimsup_{\mathsf{P} \to \infty} \frac{\sum_{\mathcal{T} \in [\mathsf{K}]^r, j \in [\mathsf{K}]\setminus\mathcal{T}} R_{\mathcal{T}}^j(\mathsf{P})}{\frac{1}{2} \log \mathsf{P}}, \tag{10}$$

where the supremum is over sequences $\{(R_{\mathcal{T}}^j(\mathsf{P}): \mathcal{T} \in [\mathsf{K}]^r, j \in [\mathsf{K}]\setminus\mathcal{T})\}_{\mathsf{P}>0}$ so that for each $\mathsf{P} > 0$ each tuple $(R_{\mathcal{T}}^j(\mathsf{P}): \mathcal{T} \in [\mathsf{K}]^r, j \in [\mathsf{K}]\setminus\mathcal{T})$ is achievable under power P.

Lemma 1. *For any* $\mathsf{r} \in [\mathsf{K}]$*:*

$$\Delta(\mathsf{r}) \geq \left(1 - \frac{\mathsf{r}}{\mathsf{K}}\right) \frac{1}{\text{Sum-DoF}(\mathsf{r})}. \tag{11}$$

Proof: See Appendix A, but the idea is well known and also used in [11] and [12]. ∎

III. NEW IA SCHEME

In view of Lemma 1, in this section we present a scheme achieving a high Sum-DoF(r) over the wireless network. We start with an example.

A. Example 2: $\mathsf{K} = 4, \mathsf{r} = 2$

Consider $\mathsf{r} = 2$ and $\mathsf{K} = 4$ nodes. In this case our scheme transmits 18 different messages depicted in (12). Here, Message $M_{k,\mathcal{T}}^j$ is a message that is known by the set of nodes $\mathcal{T}$ and intended to Node $j \notin \mathcal{T}$. (Since we consider r-fold cooperation, we have $|\mathcal{T}| = 2$.)

Message $M_{k,\mathcal{T}}^j$ is only transmitted by a single Node $k \in \mathcal{T}$. The remaining nodes in $\mathcal{T}\backslash\{k\}$ only exploit their knowledge of $M_{k,\mathcal{T}}^j$ to cancel the transmission from their receive signal.

Notice that for certain sets $\mathcal{T}$ and receive nodes $j \notin \mathcal{T}$ our scheme sends two messages to the same node j: $M_{k_1,\mathcal{T}}^j$ and $M_{k_2,\mathcal{T}}^j$ for $k_1 \neq k_2$. (In (12) the two messages $M_{1,\{1,4\}}^2$ and $M_{4,\{1,4\}}^2$ for example have this form.) These messages $M_{k_1,\mathcal{T}}^j$ and $M_{k_2,\mathcal{T}}^j$ actually represent two independent submessages of Message $M_{\mathcal{T}}^j$ as we defined it in Section II-B. For the sets $\mathcal{T}$ and Nodes $j \notin \mathcal{T}$ for which there exists only a single Message $M_{k,\mathcal{T}}^j$, this message is really the message $M_{\mathcal{T}}^j$ as we defined it in Section II-B. Since our interest is on the Sum-DoF, distinction between submessages and messages is not relevant.

In our scheme, to Node 1 we send messages

$$M_{2,\{2,3\}}^1, \ M_{3,\{2,3\}}^1, \ M_{2,\{2,4\}}^1, \ M_{3,\{3,4\}}^1; \qquad (12a)$$

to Node 2 we send messages

$$M_{1,\{1,3\}}^2, \ M_{3,\{1,3\}}^2, \ M_{1,\{1,4\}}^2, \ M_{4,\{1,4\}}^2$$
$$M_{3,\{3,4\}}^2 \ M_{4,\{3,4\}}^2; \qquad (12b)$$

to Node 3 we send messages

$$M_{1,\{1,2\}}^3, \ M_{2,\{1,2\}}^3, \ M_{1,\{1,4\}}^3, \ M_{4,\{1,4\}}^3,$$
$$M_{2,\{2,4\}}^3, \ M_{4,\{2,4\}}^3; \qquad (12c)$$

and to Node 4 we send messages

$$M_{1,\{1,2\}}^4, \ M_{2,\{1,2\}}^4, \ M_{1,\{1,3\}}^4, \ M_{3,\{1,3\}}^4,$$
$$M_{2,\{2,3\}}^4, \ M_{3,\{2,3\}}^4. \qquad (12d)$$

Node $\mathsf{K} = 4$ does not send any message to the first Node 1. (This omission allows to reuse some of the precoding matrices and achieve an improved Sum-DoF.) We encode each message $M_{k,\mathcal{T}}^j$ into a Gaussian codeword $\mathbf{b}_{k,\mathcal{T}}^j$ and use IA with three precoding matrices $\mathbf{U}_{\{2,3\}}$, $\mathbf{U}_{\{2,4\}}$, and $\mathbf{U}_{\{3,4\}}$. Matrix $\mathbf{U}_{\{2,3\}}$ is used to send codewords

$$\mathbf{b}_{1,\{1,3\}}^2, \ \mathbf{b}_{1,\{1,2\}}^3, \ \mathbf{b}_{4,\{3,4\}}^2, \ \mathbf{b}_{4,\{2,4\}}^3, \qquad (13)$$
$$\mathbf{b}_{2,\{2,3\}}^1, \ \mathbf{b}_{3,\{2,3\}}^1, \ \mathbf{b}_{3,\{1,3\}}^2, \ \mathbf{b}_{2,\{1,2\}}^3, \qquad (14)$$

matrix $\mathbf{U}_{\{2,4\}}$ for codewords

$$\mathbf{b}_{1,\{1,4\}}^2, \ \mathbf{b}_{1,\{1,2\}}^4, \ \mathbf{b}_{3,\{3,4\}}^2, \ \mathbf{b}_{3,\{2,3\}}^4, \qquad (15)$$
$$\mathbf{b}_{2,\{2,4\}}^1, \ \mathbf{b}_{2,\{1,2\}}^4, \ \mathbf{b}_{4,\{1,4\}}^2, \qquad (16)$$

and matrix $\mathbf{U}_{\{3,4\}}$ for codewords

$$\mathbf{b}_{1,\{1,4\}}^3, \ \mathbf{b}_{1,\{1,3\}}^4, \ \mathbf{b}_{2,\{2,4\}}^3, \ \mathbf{b}_{2,\{2,3\}}^4, \qquad (17)$$
$$\mathbf{b}_{3,\{3,4\}}^1, \ \mathbf{b}_{4,\{1,4\}}^3, \ \mathbf{b}_{3,\{1,3\}}^4. \qquad (18)$$

TABLE I
MESSAGES $M_{k,\mathcal{T}}^j$ PRECODED BY THE THREE PRECODING MATRICES $\mathbf{U}_{\{2,3\}}$, $\mathbf{U}_{\{2,4\}}$, AND $\mathbf{U}_{\{3,4\}}$.

$\mathcal{T} \backslash j$	1	2	3	4
$\{1,2\}$	x	x	$\mathbf{U}_{\{2,3\}}$	$\mathbf{U}_{\{2,4\}}$
$\{1,3\}$	x	$\mathbf{U}_{\{2,3\}}$	x	$\mathbf{U}_{\{3,4\}}$
$\{1,4\}$	x	$\mathbf{U}_{\{2,4\}}$	$\mathbf{U}_{\{3,4\}}$	x
$\{2,3\}$	$\mathbf{U}_{\{2,3\}}$	x	x	$\mathbf{U}_{\{2,4\}},$ $\mathbf{U}_{\{3,4\}}$
$\{2,4\}$	$\mathbf{U}_{\{2,4\}}$	x	$\mathbf{U}_{\{2,3\}},$ $\mathbf{U}_{\{3,4\}}$	x
$\{3,4\}$	$\mathbf{U}_{\{3,4\}}$	$\mathbf{U}_{\{2,3\}}, \mathbf{U}_{\{2,4\}}$	x	x

Remark 1. *The choice of precoding matrices is inspired by [14] where Message $M_{k,\mathcal{T}}^j$ is precoded by the matrix $\mathbf{U}_{\mathcal{R}}$ for $\mathcal{R} = \mathcal{T}\backslash\{j\} \cup \{k\}$. The idea behind the choice of precoding matrices in [14] is that any node in $\mathcal{R}$ is either interested in learning Message $M_{k,\mathcal{T}}^j$ or it can compute it itself and remove the interference from its receive signal. A given node j thus only experiences interference from precoding matrices $\mathbf{U}_{\mathcal{R}}$ for which $j \notin \mathcal{R}$. In our IA scheme, we omit precoding matrices $\mathbf{U}_{\mathcal{R}'}$ for sets $\mathcal{R}'$ containing index 1, and instead use the precoding matrix $\mathbf{U}_{\mathcal{R}}$ also to send*

$$\mathbf{b}_{k,\mathcal{R}}^1, \qquad \mathbf{b}_{k,\mathcal{R}\cup\{1\}\backslash\{j\}}^j, \qquad \forall j,k \in \mathcal{R}, \ j \neq k, \qquad (19)$$

see the codewords indicated in (14), (16), (18).

We illustrate our assignment of the precoding matrices also in Table I. The entries in column 1 or in rows $\{1,2\}, \{1,3\}, \{1,4\}$ correspond to two submessages $M_{k_1,\mathcal{T}}^j$ and $M_{k_2,\mathcal{T}}^j$, where k_1 and k_2 denote the two entries in $\mathcal{T}$. For all other entries in Table I not equal to "x", we have only one message per precoding matrix, see (13), (15), and (17).

During the shuffling phase, Nodes 1–4 send the following signals. Node 1 sends:

$$\mathbf{X}_1 = \mathbf{U}_{\{2,3\}}\left(\mathbf{b}_{1,\{1,3\}}^2 + \mathbf{b}_{1,\{1,2\}}^3\right)$$
$$+ \mathbf{U}_{\{2,4\}}\left(\mathbf{b}_{1,\{1,4\}}^2 + \mathbf{b}_{1,\{1,2\}}^4\right)$$
$$+ \mathbf{U}_{\{3,4\}}\left(\mathbf{b}_{1,\{1,4\}}^3 + \mathbf{b}_{1,\{1,3\}}^4\right). \qquad (20)$$

Node 2 sends:

$$\mathbf{X}_2 = \mathbf{U}_{\{2,3\}}\left(\mathbf{b}_{2,\{2,3\}}^1 + \mathbf{b}_{2,\{1,2\}}^3\right)$$
$$+ \mathbf{U}_{\{2,4\}}\left(\mathbf{b}_{2,\{2,4\}}^1 + \mathbf{b}_{2,\{1,2\}}^4\right)$$
$$+ \mathbf{U}_{\{3,4\}}\left(\mathbf{b}_{2,\{2,4\}}^3 + \mathbf{b}_{2,\{2,3\}}^4\right). \qquad (21)$$

Node 3 sends:

$$\mathbf{X}_3 = \mathbf{U}_{\{2,3\}}\left(\mathbf{b}_{3,\{2,3\}}^1 + \mathbf{b}_{3,\{1,3\}}^2\right)$$
$$+ \mathbf{U}_{\{2,4\}}\left(\mathbf{b}_{3,\{3,4\}}^2 + \mathbf{b}_{3,\{2,3\}}^4\right)$$
$$+ \mathbf{U}_{\{3,4\}}\left(\mathbf{b}_{3,\{3,4\}}^1 + \mathbf{b}_{3,\{1,3\}}^4\right). \qquad (22)$$

Node 4 sends:

$$\mathbf{X}_4 = \mathbf{U}_{\{2,3\}} \left(\mathbf{b}^2_{4,\{3,4\}} + \mathbf{b}^3_{4,\{2,4\}} \right)$$
$$+ \mathbf{U}_{\{2,4\}} \mathbf{b}^2_{4,\{1,4\}} + \mathbf{U}_{\{3,4\}} \mathbf{b}^3_{4,\{1,4\}}. \qquad (23)$$

Eeach node subtracts all the interference of the signals that it can compute itself. For example, Node 2 thus constructs:

$$\mathbf{Y}'_2 = \underbrace{\mathbf{H}_{2,1} \mathbf{U}_{\{2,3\}} \mathbf{b}^2_{1,\{1,3\}} + \mathbf{H}_{2,1} \mathbf{U}_{\{2,4\}} \mathbf{b}^2_{1,\{1,4\}}}_{\text{desired signal}}$$
$$+ \underbrace{\mathbf{H}_{2,3} \mathbf{U}_{\{2,3\}} \mathbf{b}^2_{3,\{1,3\}} + \mathbf{H}_{2,3} \mathbf{U}_{\{2,4\}} \mathbf{b}^2_{3,\{3,4\}}}_{\text{desired signal}}$$
$$+ \underbrace{\mathbf{H}_{2,4} \mathbf{U}_{\{2,3\}} \mathbf{b}^2_{4,\{3,4\}} + \mathbf{H}_{2,4} \mathbf{U}_{\{2,4\}} \mathbf{b}^2_{4,\{1,4\}}}_{\text{desired signal}}$$
$$+ \mathbf{H}_{2,1} \mathbf{U}_{\{3,4\}} \left(\mathbf{b}^3_{1,\{1,4\}} + \mathbf{b}^4_{1,\{1,3\}} \right)$$
$$+ \mathbf{H}_{2,3} \mathbf{U}_{\{3,4\}} \left(\mathbf{b}^1_{3,\{3,4\}} + \mathbf{b}^4_{3,\{1,3\}} \right)$$
$$+ \mathbf{H}_{2,4} \mathbf{U}_{\{3,4\}} \mathbf{b}^3_{4,\{2,4\}} + \mathbf{Z}_2. \qquad (24)$$

Node 2's desired signals are all precoded by precoding matrices $\mathbf{U}_{\{2,3\}}$ and $\mathbf{U}_{\{2,4\}}$, while all interference signals are precoded by matrix $\mathbf{U}_{\{3,4\}}$. Similar observations hold for Nodes 3 and 4. Desired and interference signals at Node 1 instead are precoded by all three precoding matrices.

The IA matrices $\mathbf{U}_{\{2,3\}}$, $\mathbf{U}_{\{2,4\}}$, and $\mathbf{U}_{\{3,4\}}$ are constructed based on the IA idea in [15] taking into account the channel matrices that premultiply the IA matrices in the interference signals of $\mathbf{Y}'_1, \mathbf{Y}'_2, \mathbf{Y}'_3, \mathbf{Y}'_4$. Specifically:

$$\mathbf{U}_{\mathcal{R}} \triangleq \left[\prod_{\mathbf{H} \in \mathcal{H}_{\mathcal{R}}} \mathbf{H}^{\alpha_{\mathcal{R}}, \mathbf{H}} \cdot \boldsymbol{\Xi}_{\mathcal{R}} : \ \forall \boldsymbol{\alpha}_{\mathcal{R}} \in [\eta]^4 \right], \qquad (25)$$

where each column of the matrix is constructed using a different exponent-vector $\boldsymbol{\alpha}_{\mathcal{R}} = (\alpha_{\mathcal{R},\mathbf{H}} : \mathbf{H} \in \mathcal{H}_{\mathcal{R}}) \in [\eta]^4$; η is a large number depending on the blocklength T that tends to ∞ with T; $\boldsymbol{\Xi}_{\mathcal{R}}$ are i.i.d. random vectors drawn according to a continuous distribution, and

$$\mathcal{H}_{\{2,3\}} \triangleq \{ \mathbf{H}_{1,4}, \ \mathbf{H}_{4,1}, \ \mathbf{H}_{4,2}, \ \mathbf{H}_{4,3} \}, \qquad (26)$$
$$\mathcal{H}_{\{2,4\}} \triangleq \{ \mathbf{H}_{1,3}, \ \mathbf{H}_{3,1}, \ \mathbf{H}_{3,2}, \ \mathbf{H}_{3,4} \}, \qquad (27)$$
$$\mathcal{H}_{\{3,4\}} \triangleq \{ \mathbf{H}_{1,2}, \ \mathbf{H}_{2,1}, \ \mathbf{H}_{2,3}, \ \mathbf{H}_{2,4} \}. \qquad (28)$$

By this choice of the precoding matrices, all interference signals at a Node 2 will lie in the columnspace of the matrix

$$\mathbf{W}_{\{3,4\}} \triangleq \left[\prod_{\mathbf{H} \in \mathcal{H}_{\{3,4\}}} \mathbf{H}^{\alpha_{\mathcal{R}}, \mathbf{H}} \cdot \boldsymbol{\Xi}_{\{3,4\}} : \ \forall \boldsymbol{\alpha}_{\mathcal{R}} \in [\eta + 1]^4 \right],$$

while the desired signals will be separable from each other and from this interference space.

Since $\mathbf{Y}'_2$ has 6 signal spaces for the 6 desired signals and only a single interference space, our scheme achieves DoF $6/7$ to Node 2. Similar considerations hold for Nodes 3 and 4. Node 1 has 3 interference spaces (one for each IA precoding matrix) and only 4 signal spaces, and we thus achieve DoF

$4/7$ to Node 1. The SumDoF achieved by the scheme is thus: Sum-DoF $= 22/7$.

For comparison, notice that the IA scheme in [14] uses the additional precoding matrices $\mathbf{U}_{\{1,2\}}$, $\mathbf{U}_{\{1,3\}}$, and $\mathbf{U}_{\{1,4\}}$. The interference space at Node 2 then consists of $\mathbf{W}_{\{3,4\}}$ and additionally of the similar matrices $\mathbf{W}_{\{1,2\}}$ and $\mathbf{W}_{\{1,4\}}$. Thus, only DoF $6/(6+3) = 2/3$ is achievable to Node 2, as well as to all other nodes, leading to the reduced Sum-DoF $12/3$.

B. The General Scheme

Fix $\eta \in \mathbb{Z}^+$ (which we shall let tend to ∞) and let

$$\Gamma \triangleq \mathsf{K} \cdot (\mathsf{K} - \mathsf{r} - 1) \qquad (29)$$
$$\mathsf{T} \triangleq (\mathsf{K} - 2) \cdot \binom{\mathsf{K} - 2}{\mathsf{r} - 1} \cdot \eta^{\Gamma} + \binom{\mathsf{K} - 1}{\mathsf{r}} \cdot (\eta + 1)^{\Gamma}. \qquad (30)$$

We send the following messages to any Node $j \in [\mathsf{K}] \backslash \{1\}$:

$$\left\{ M^j_{k,\mathcal{T}} : \mathcal{T} \in [[\mathsf{K}] \backslash \{j\}]^{\mathsf{r}}, \ k \in \mathcal{T} \right\} \qquad (31)$$

and to Node 1 we send messages

$$\left\{ M^1_{k,\mathcal{T}} : \mathcal{T} \in [[\mathsf{K}] \backslash \{1\}]^{\mathsf{r}}, \ k \in \mathcal{T} \backslash \{\mathsf{K}\} \right\}. \qquad (32)$$

Thus, as in the examples of the previous section, the last node K does not send any message to the Node 1.

For each message, construct a Gaussian codebook of power $\mathsf{P} / \binom{\mathsf{K} - 1}{\mathsf{r}}$ and length η^{Γ} to encode each Message $M^j_{k,\mathcal{T}}$ into a codeword $\mathbf{b}^j_{k,\mathcal{T}}$. As in the previous sections, we shall use a linear precoding scheme, and thus Node $i \in [\mathsf{K}]$ can mitigate the interference caused by the codewords

$$\left\{ \mathbf{b}^j_{k,\mathcal{T}} \right\}_{\forall \mathcal{T} : i \in \mathcal{T}}. \qquad (33)$$

Thus, for each set $\mathcal{R} \in [[\mathsf{K}]]^{\mathsf{r}}$, without causing non-desired interference to nodes in $\mathcal{R}$, we can use the same precoding matrix $\mathbf{U}_{\mathcal{R}}$ for all the codewords:

$$\left\{ \mathbf{b}^j_{k,\mathcal{R} \cup \{k\} \backslash \{j\}} \right\}_{\substack{k \in [\mathsf{K}] \backslash \mathcal{R} \\ j \in \mathcal{R}}}. \qquad (34)$$

This idea was already used in the related works [13], [14]. In contrast to these previous works, here we do not introduce the precoding matrices $\mathbf{U}_{\mathcal{R}}$ for sets $\mathcal{R}$ containing 1 and instead we use matrix $\mathbf{U}_{\mathcal{R}}$, for $1 \notin \mathcal{R}$, also to precode the codewords

$$\left\{ \mathbf{b}^j_{k,\mathcal{R} \cup \{1\} \backslash \{j\}} \right\}_{j,k \in \mathcal{R}, j \neq k} \quad \cup \quad \left\{ \mathbf{b}^j_{1,\mathcal{R} \cup \{1\} \backslash \{j\}} \right\}_{j \in \mathcal{R}}, \qquad (35)$$

and

$$\left\{ \mathbf{b}^1_{k,\mathcal{R}} \right\}_{k \in \mathcal{R} \backslash \{\mathsf{K}\}}. \qquad (36)$$

All non-intended nodes in $\mathcal{R}$ can subtract these interferences from their receive signals because they know the codewords. This trick allows us to reduce the dimension of the interference space and thus improve performance.

Table II illustrates which codesymbols $\mathbf{b}^j_{k,\mathcal{T}}$ are premultiplied by the precoding matrix $\mathbf{U}_{\{2,3,4\}}$, when $\mathsf{r} = 3$ and $\mathsf{K} \geq 5$. The entries in rows $\mathcal{T}$ containing index 1 correspond to $\mathsf{r} = 3$ different codewords $\mathbf{b}^j_{k,\mathcal{T}}$, one for each $k \in \mathcal{T}$, see (35). Similarly, the entry in column-1 and row $\{2, 3, 4\}$ corresponds to the r codewords $\mathbf{b}^1_{k,\{2,3,4\}}$, for each $k \in \{2, 3, 4\}$. Any other

$\mathcal{T} \setminus j$	1	2	3	4
$\{1,2,3\}$	x	x	x	$\mathbf{U}_{\{2,3,4\}}$
$\{1,2,4\}$	x	x	$\mathbf{U}_{\{2,3,4\}}$	x
$\{1,3,4\}$	x	$\mathbf{U}_{\{2,3,4\}}$	x	x
$\{2,3,4\}$	$\mathbf{U}_{\{2,3,4\}}$	x	x	x
$\{2,3,5\}$	o	x	x	$\mathbf{U}_{\{2,3,4\}}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\{2,3,\mathsf{K}\}$	o	x	x	$\mathbf{U}_{\{2,3,4\}}$
$\{2,3,5\}$	o	x	$\mathbf{U}_{\{2,3,4\}}$	x
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\{2,3,\mathsf{K}\}$	o	x	$\mathbf{U}_{\{2,3,4\}}$	x
$\{3,4,5\}$	o	$\mathbf{U}_{\{2,3,4\}}$	x	x
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\{3,4,\mathsf{K}\}$	o	$\mathbf{U}_{\{2,3,4\}}$	x	x

entry of the table showing $\mathbf{U}_{\{2,3,4\}}$ corresponds to a single codeword $\mathbf{b}_{k,\mathcal{T}}^{j}$, where k is the single element in $\mathcal{T} \setminus \{2,3,4\}$. Similar tables can be drawn for all pairs $(k_1, k_2) \in [\mathsf{K}]$, where recall that node K does not send any information to Node 1.

Encoding: Define the T-length vector of channel inputs $\mathbf{X}_k \triangleq (X_k(1), \ldots, X_k(\mathsf{T}))^T$ for each Node k and set:

$$\mathbf{X}_1 = \sum_{\mathcal{R} \in [[\mathsf{K}] \setminus \{1\}]^r} \sum_{j \in \mathcal{R}} \mathbf{U}_{\mathcal{R}} \mathbf{b}_{1,\mathcal{R} \cup \{1\} \setminus \{j\}}^{j}, \tag{37}$$

$$\mathbf{X}_k = \sum_{\mathcal{R} \in [[\mathsf{K}] \setminus \{1,k\}]^r} \sum_{j \in \mathcal{R}} \mathbf{U}_{\mathcal{R}} \mathbf{b}_{k,\mathcal{R} \cup \{k\} \setminus \{j\}}^{j}$$
$$+ \sum_{\substack{\mathcal{R} \in [[\mathsf{K}] \setminus \{k\}]^r: \\ 1 \in \mathcal{R}}} \sum_{j \in \mathcal{R}} \mathbf{U}_{\mathcal{R} \cup \{k\} \setminus \{1\}} \mathbf{b}_{k,\mathcal{R} \cup \{k\} \setminus \{j\}}^{j}, \\ k \in [\mathsf{K}-1] \setminus \{1\}, \tag{38}$$

$$\mathbf{X}_{\mathsf{K}} = \sum_{\mathcal{R} \in [[\mathsf{K}-1] \setminus \{1\}]^r} \sum_{j \in \mathcal{R}} \mathbf{U}_{\mathcal{R}} \mathbf{b}_{\mathsf{K},\mathcal{R} \cup \{\mathsf{K}\} \setminus \{j\}}^{j}$$
$$+ \sum_{\substack{\mathcal{R} \in [[\mathsf{K}-1]]^r: \\ 1 \in \mathcal{R}}} \sum_{j \in \mathcal{R} \setminus \{1\}} \mathbf{U}_{\mathcal{R} \cup \{\mathsf{K}\} \setminus \{1\}} \mathbf{b}_{\mathsf{K},\mathcal{R} \cup \{\mathsf{K}\} \setminus \{j\}}^{j}, \tag{39}$$

where we shortly describe matrices $\{\mathbf{U}_{\mathcal{R}}\}_{\mathcal{R} \in [[\mathsf{K}] \setminus \{1\}]^r}$.

Decoding: After receiving the respective sequence of T channel outputs $\mathbf{Y}_j \triangleq (Y_{j,1}, \ldots, Y_{j,\mathsf{T}})$, for $j \in [\mathsf{K}]$, each node removes the influence of the codewords corresponding to the messages that it can compute itself. The nodes' "cleaned" signals can then be written as:

$$\mathbf{Y}_1' = \underbrace{\sum_{\substack{\mathcal{R} \in [[\mathsf{K}]]^r: \\ 1 \in \mathcal{R}}} \sum_{k \in [\mathsf{K}-1] \setminus \mathcal{R}} \mathbf{H}_{1,k} \mathbf{U}_{\mathcal{R} \cup \{k\} \setminus \{1\}} \mathbf{b}_{k,\mathcal{R} \cup \{k\} \setminus \{1\}}^{1}}_{\text{desired signal}}$$
$$+ \sum_{\mathcal{R} \in [[\mathsf{K}] \setminus \{1\}]^r} \sum_{k \in [\mathsf{K}] \setminus \mathcal{R}} \mathbf{H}_{1,k} \mathbf{U}_{\mathcal{R}} \mathbf{v}_{\mathcal{R},k} + \mathbf{Z}_1, \tag{40a}$$

$$\mathbf{Y}_j' = \underbrace{\sum_{\substack{\mathcal{R} \in [[\mathsf{K}] \setminus \{1\}]^r: \\ j \in \mathcal{R}}} \sum_{k \in [\mathsf{K}] \setminus \mathcal{R}} \mathbf{H}_{j,k} \mathbf{U}_{\mathcal{R}} \mathbf{b}_{k,\mathcal{R} \cup \{k\} \setminus \{j\}}^{j}}_{\text{desired signal}}$$

$$+ \underbrace{\sum_{\substack{\mathcal{R} \in [[\mathsf{K}]]^r: \\ 1,j \in \mathcal{R}}} \sum_{k \in [\mathsf{K}] \setminus \mathcal{R}} \mathbf{H}_{j,k} \mathbf{U}_{\mathcal{R} \cup \{k\} \setminus \{1\}} \mathbf{b}_{k,\mathcal{R}}^{j}}_{\text{desired signal}}$$
$$+ \sum_{\substack{\mathcal{R} \in [[\mathsf{K}] \setminus \{1\}]^r: \\ j \notin \mathcal{R}}} \sum_{\substack{k \in [\mathsf{K}] \setminus \mathcal{R}: \\ k \neq j}} \mathbf{H}_{j,k} \mathbf{U}_{\mathcal{R}} \mathbf{v}_{\mathcal{R},k}$$
$$+ \sum_{\substack{\mathcal{R} \in [[\mathsf{K}]]^r: \\ 1 \in \mathcal{R}, j \notin \mathcal{R}}} \sum_{k \in [\mathsf{K}] \setminus \mathcal{R}} \mathbf{H}_{j,k} \mathbf{U}_{\mathcal{R} \cup \{k\} \setminus \{1\}} \mathbf{v}_{\mathcal{R},k} + \mathbf{Z}_j, \\ j \in [\mathsf{K}] \setminus \{1\}, \tag{40b}$$

where for ease of notation we defined for Nodes $k \in [\mathsf{K}-1]$:

$$\mathbf{v}_{\mathcal{R},k} \triangleq \sum_{j \in \mathcal{R}} \mathbf{b}_{k,\mathcal{R} \cup \{k\} \setminus \{j\}}^{j}, \qquad \forall \mathcal{R} \in [[\mathsf{K}] \setminus \{k\}]^r, \tag{41}$$

and for the last Node K, since its signal to Node 1 is absent:

$$\mathbf{v}_{\mathcal{R},\mathsf{K}} \triangleq \sum_{j \in \mathcal{R} \setminus \{1\}} \mathbf{b}_{k,\mathcal{R} \cup \{k\} \setminus \{j\}}^{j}, \qquad \forall \mathcal{R} \in [[\mathsf{K}-1]]^r. \tag{42}$$

Each Node j zero-forces the non-desired interference terms of its "cleaned" signal and decodes its intended messages.

Choice of IA Matrices $\{\mathbf{U}_{\mathcal{R}}\}$: Inspired by the IA scheme in [15], we choose each $\mathsf{T} \times \eta^{\Gamma}$ precoding matrix $\mathbf{U}_{\mathcal{R}}$ so that its column-span includes all power products (powers 1 to η) of the channel matrices $\mathbf{H}_{j,k}$ that premultiply $\mathbf{U}_{\mathcal{R}}$ in (40) in the non-desired interference terms. Thus, $\mathcal{R} \in [[\mathsf{K}] \setminus \{1\}]^r$:

$$\mathbf{U}_{\mathcal{R}} \triangleq \left[\prod_{\mathbf{H} \in \mathcal{H}_{\mathcal{R}}} \mathbf{H}^{\alpha_{\mathcal{R},\mathbf{H}}} \cdot \boldsymbol{\Xi}_{\mathcal{R}} : \ \forall \boldsymbol{\alpha}_{\mathcal{R}} \in [\eta]^{\Gamma} \right], \tag{43}$$

where $\{\boldsymbol{\Xi}_{\mathcal{R}}\}_{\mathcal{R} \in [[\mathsf{K}] \setminus \{1\}]^r}$ are i.i.d. random vectors independent of all channel matrices, noises, and messages,

$$\mathcal{H}_{\mathcal{R}} \triangleq \big\{ \mathbf{H}_{j,k} : \ j \in [\mathsf{K}] \setminus \mathcal{R}, \ k \in [\mathsf{K}] \setminus \{j\} \big\} \setminus \big\{ \mathbf{H}_{1,k} : k \in \mathcal{R} \big\},$$

and $\boldsymbol{\alpha}_{\mathcal{R}} \triangleq (\alpha_{\mathcal{R},\mathbf{H}} : \mathbf{H} \in \mathcal{H}_{\mathcal{R}})$. Notice that $|\mathcal{H}_{\mathcal{R}}| = \Gamma$ for any $\mathcal{R} \in [[\mathsf{K}] \setminus \{1\}]^r$.

Performance Analysis: See Appendix B.

IV. NEW BOUNDS ON THE NDT

Define for any integer value $r \in [\mathsf{K}]$:

$$\Delta_{\mathsf{Ub}}(r) \triangleq \begin{cases} \left(1 - \frac{r}{\mathsf{K}}\right) \cdot \frac{r(\mathsf{K}-1) + \mathsf{K} - r - 1}{r(\mathsf{K}-1)^2 + r(\mathsf{K}-2)} & \text{if } r < \mathsf{K}/2 \\ \frac{1}{\mathsf{K}} \left(1 - \frac{r}{\mathsf{K}}\right) & \text{if } r \geq \mathsf{K}/2 \end{cases}. \tag{44}$$

Also, let

$$\Delta_{\mathsf{Lb}}(r) \triangleq \begin{cases} \frac{1}{\mathsf{K}} \left(2 - \frac{3}{\mathsf{K}}\right) & \text{if } r = 1, \\ \frac{1}{\mathsf{K}} \left(1 - \frac{r}{\mathsf{K}} + \max_{t \in [\lfloor \mathsf{K}/2 \rfloor]} \mathrm{lowc}\left(C_t(r)\right)\right) & \text{if } r \in (1,2), \\ \frac{1}{\mathsf{K}} \left(1 - \frac{r}{\mathsf{K}} + \mathrm{lowc}\left(C_{\lfloor \mathsf{K}/2 \rfloor}(r)\right)\right) & \text{if } r \in [2,\mathsf{K}], \end{cases} \tag{45}$$

where for any $t \in [\lfloor \mathsf{K}/2 \rfloor]$:

$$C_t(i) = \begin{cases} \frac{\binom{\mathsf{K}-i}{t-i}}{\binom{\mathsf{K}}{t} \cdot t} \cdot (\mathsf{K}-t-i), & \text{if } i \in [t], \\ 0, & \text{if } i \in [\mathsf{K}] \setminus [t] \end{cases} \tag{46}$$

5

and for any function f, lowc $(f(\ell))$ denotes the lower convex enveloppe of $\{(\ell, f(\ell))\}$.

Theorem 1. *The NDT-computation tradeoff $\Delta^*(\mathsf{r})$ is upper- and lower-bounded as:*

$$\Delta_{\text{Lb}}(\mathsf{r}) \leq \Delta^*(\mathsf{r}) \leq \text{lowc}\left(\Delta_{\text{Ub}}(\mathsf{r})\right). \tag{47}$$

Proof: For integers $\mathsf{r} \geq \mathsf{K}/2$ achievability of $\Delta_{\text{Ub}}(\mathsf{r})$ is proved in [11]. For integers $\mathsf{r} < \mathsf{K}/2$ achievability of $\Delta_{\text{Ub}}(\mathsf{r})$ holds by Lemma 1 and the scheme in Section III. Achievability of the lower convex enveloppe follows by simple time- and memory-sharing strategies. The lower bound can be proved using MAC-type arguments, see Appendix C. ∎

Remark 2. *The upper bound is convex and piece-wise constant. The lower bound is piecewise constant with segments spanning the intervals $[i, i+1]$, for $i = 2, \ldots, \mathsf{K} - 1$. On the interval $[1, 2)$, the lower bound is constant over smaller sub-intervals only but not over the entire segment.*

Corollary 1. *For all $\mathsf{r} \geq \lceil \mathsf{K}/2 \rceil$, the linear interference cancelation scheme in [11] achieves the NDT, which equals*

$$\Delta^*(\mathsf{r}) = \left(1 - \frac{\mathsf{r}}{\mathsf{K}}\right) \cdot \frac{1}{\mathsf{K}}. \tag{48}$$

Proof: For $\mathsf{r} \geq \lceil \mathsf{K}/2 \rceil$ the upper bound lowc $(\Delta_{\text{Ub}}(\mathsf{r}))$ is equal to the lower bound $\Delta_{\text{Lb}}(\mathsf{r})$ because $C_{\lfloor \mathsf{K}/2 \rfloor}(i) = 0$ for all $i \geq \lceil \mathsf{K}/2 \rceil$. ∎

Remark 3. *By [11], $\Delta^*(\mathsf{r})$ in (48) is achieved with beamforming, zero-forcing, and side-information cancellation. By Corollaries 1 and 2, these simple strategies are sufficient to achieve $\Delta^*(\mathsf{r})$ when $\mathsf{r} \geq \lceil \mathsf{K}/2 \rceil$ but not when $\mathsf{r} < \lceil \frac{\mathsf{K}-1}{2} \rceil$.*

We compare the upper bound in Theorem 1 to the bounds in [11] and [12]. The upper bound in [11] is given as follows:

$$\Delta^*(\mathsf{r}) \leq \Delta_{\text{UB-BF}}(\mathsf{r}) \triangleq \text{lowc}\left\{\left(\mathsf{r}, \frac{1 - \mathsf{r}/\mathsf{K}}{\min(\mathsf{K}, 2\mathsf{r})}\right) : \mathsf{r} \in [\mathsf{K}]\right\}. \tag{49}$$

The upper bound in [12] has the form:

$$\Delta^*(\mathsf{r}) \leq \Delta_{\text{Ub-Groups}}(\mathsf{r}) \triangleq$$
$$\text{lowc}\left((\mathsf{K}, 0) \cup \left\{\left(\mathsf{r}, \frac{1 - \mathsf{r}/\mathsf{K}}{\text{Sum-DoF}_{\text{Lb}}(\mathsf{r})}\right) : 1 \leq \mathsf{r} < \mathsf{K}, \mathsf{r} | \mathsf{K}\right\}\right), \tag{50}$$

where

$$\text{Sum-DoF}_{\text{Lb}}(\mathsf{r}) \triangleq \begin{cases} 2\mathsf{r} & \text{if } \mathsf{K}/\mathsf{r} \in \{2, 3\}, \\ \frac{\mathsf{K}(\mathsf{K}-\mathsf{r})-\mathsf{r}^2}{2\mathsf{K}-3\mathsf{r}} & \text{if } \mathsf{K}/\mathsf{r} \geq 4. \end{cases} \tag{51}$$

Notice that $\Delta_{\text{Ub}}(1) = \Delta_{\text{Ub-Groups}}(1)$.

Corollary 2. *For all $1 < \mathsf{r} < \lceil \frac{\mathsf{K}}{2} \rceil$:*

$$\Delta^*(\mathsf{r}) \leq \text{lowc}\left(\Delta_{\text{Ub}}(\mathsf{r})\right) < \Delta_{\text{Ub-Groups}}(\mathsf{r}), \tag{52}$$

and for all $1 \leq \mathsf{r} < \lceil \frac{\mathsf{K}-1}{2} \rceil$:

$$\Delta^*(\mathsf{r}) \leq \text{lowc}\left(\Delta_{\text{Ub}}(\mathsf{r})\right) < \Delta_{\text{Ub-ZF}}(\mathsf{r}). \tag{53}$$

Fig. 1 and 2 compare the bounds in Theorem 1 to the previous upper bounds $\Delta_{\text{Ub-Groups}}(\mathsf{r})$ and $\Delta_{\text{UB-BF}}(\mathsf{r})$.
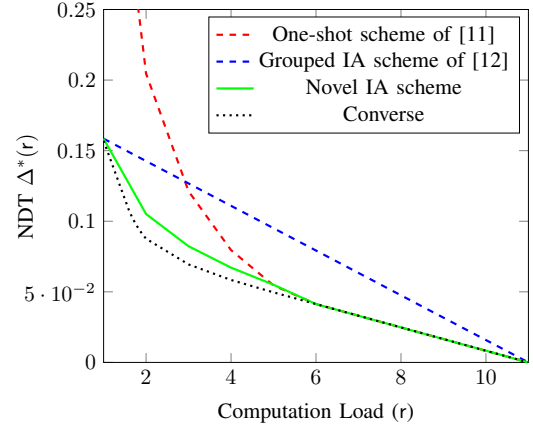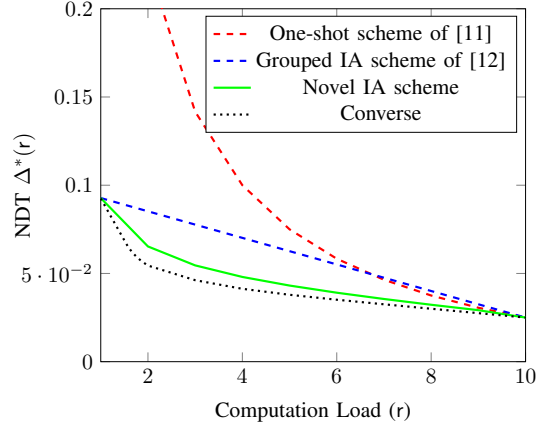

Fig. 1. Bounds on $\Delta^*(\mathsf{r})$ for $\mathsf{K} = 11$.

Fig. 2. Bounds on $\Delta^*(\mathsf{r})$ for $\mathsf{K} = 20$.

V. Conclusion

This paper presents an improved upper bound and the first information-theoretic lower bound on the NDT tradeoff of full-duplex wireless MapReduce systems. The upper bound is obtained by zero-forcing and a novel IA scheme that is tailored to the information cancellation capabilities of the nodes in a MapReduce system. As a conclusion of this work, we observe that linear beamforming, zero-forcing, and interference cancelation are optimal when each node can store at least half of the file, but they are suboptimal for smaller computation loads.

References

[1] J. S. Ng, W. Y. B. Lim, N. C. Luong, Z. Xiong, A. Asheralieva, D. Niyato, C. Leung, and C. Miao, "A comprehensive survey on coded distributed computing: Fundamentals, challenges, and networking applications," *IEEE Commun. Surveys Tuts.*, vol. 23, no. 3, pp. 1800–1837, Jun. 2021.

[2] J. Dean and S. Ghemawat, "MapReduce: simplified data processing on large clusters," *Comm. of ACM*, vol. 51, no. 1, pp. 107–113, Jan. 2008.

[3] M. Chowdhury, M. Zaharia, J. Ma, M. I. Jordan, and I. Stoica, "Managing data transfers in computer clusters with orchestra," *ACM SIGCOMM computer communication review*, vol. 41, no. 4, pp. 98–109, 2011.

[4] S. Li, M. A. Maddah-Ali, Q. Yu, and A. S. Avestimehr, "A fundamental tradeoff between computation and communication in distributed computing," *IEEE Trans. Inf. Th.*, vol. 64, no. 1, pp. 109–128, Jan. 2018.

[5] Q. Yan, S. Yang, and M. Wigger, "Storage-computation-communication tradeoff in distributed computing: Fundamental limits and complexity," *IEEE Trans. Inf. Th.*, vol. 68, no. 8, pp. 5496–5512, 2022.

[6] F. Xu, S. Shao, and M. Tao, "New results on the computation-communication tradeoff for heterogeneous coded distributed computing," *IEEE Trans. Commun.*, vol. 69, no. 4, pp. 2254–2270, Apr. 2021.

[7] S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, "A unified coding framework for distributed computing with straggling servers," in *Proc. IEEE Globecom Wkshps*, Dec. 2016, pp. 1–6.

[8] J. Zhang and O. Simeone, "Improved latency-communication trade-off for map-shuffle-reduce systems with stragglers," in *Proc. IEEE ICASSP*, May 2019, pp. 8172–8176, iSSN: 2379-190X.

[9] L. Zheng, M. Sun, Y. Luo, X. Song, C. Yang, F. Hu, and M. Yu, "Utilizing mapreduce to improve probe-car track data mining," *ISPRS International Journal of Geo-Information*, vol. 7, no. 7, 2018.

[10] Y. Khourdifi and M. Bahaj, "Hadoop and MapReduce technology as a solution for wireless body area networks in e-health," in *Proc. of ICCWCS*, New York, NY, USA, 2017.

[11] F. Li, J. Chen, and Z. Wang, "Wireless mapreduce distributed computing," *IEEE Trans. Inf. Th.*, vol. 65, no. 10, pp. 6101–6114, Oct. 2019.

[12] Y. Bi, P. Ciblat, M. Wigger, and Y. Wu, "DoF of a cooperative X-channel with an application to distributed computing," in *Proc. IEEE ISIT*, Espoo, Finland, Jun. 2022, pp. 566–571.

[13] K. Yuan and Y. Wu, "Coded wireless distributed computing via inter-ference alignment," in *Proc. IEEE ISIT*, Espoo, Finland, 2022.

[14] J. Hachem, U. Niesen, and S. N. Diggavi, "Degrees of freedom of cache-aided wireless interference networks," *IEEE Trans. Inf. Th.*, vol. 64, no. 7, pp. 5359–5380, Jul. 2018.

[15] S. Jafar and S. Shamai, "Degrees of freedom region of the MIMO X-channel," *IEEE Trans. Inf. Th.*, vol. 54, no. 1, pp. 151–170, Jan. 2008.

APPENDIX A
PROOF OF LEMMA 1

We show how to construct a distributed computing scheme achieving the NDT upper bound in (11). We shall assume a sequence (in $P > 0$) of rates $(R_\mathcal{T}^j : \mathcal{T} \in [K]^r, \ j \in [K] \backslash \mathcal{T})$ that achieves the sum-DoF Sum-DoF(r) and is completely symmetric with respect to indices j and sets $\mathcal{T}$. By the same time-sharing and relabeling arguments as described in Subsection II such a sequence must exist.

In the Map Phase we choose a regular file assignment. Partition the input files $\{W_1, \ldots, W_N\}$ into $\binom{K}{r}$ disjoint bundles and assign each bundle to a size-r subset $\mathcal{T} \in [[K]]^r$. The bundle associated to subset $\mathcal{T}$ is denoted $\mathcal{W}_\mathcal{T} \subseteq \{W_1, \ldots, W_N\}$. Notice that the proposed assignment satisfies the constraint on the computation load, because the number of files stored at Node k is:

$$\binom{K-1}{r-1} \frac{N}{\binom{K}{r}} = \frac{r}{K}N. \tag{54}$$

Each node computes all IVAs associated to its stored files.

During the Wireless Shuffle Phase, each transmit set $\mathcal{T}$ communicates to any receive node $j \notin \mathcal{T}$ all IVAs that can be calculated from bundle $\mathcal{W}_\mathcal{T}$. To this end, all nodes use the encoding and decoding functions achieving Sum-DoF(r).

Each transmit group $\mathcal{T}$ has to send $\frac{N}{\binom{K}{r}}$ IVAs to each Node $j \notin \mathcal{T}$ and in total there are $\binom{K}{r}(K-r)$ rates in the symmetric rate vector $(R_\mathcal{T}^j : \mathcal{T} \in [K]^r, \ j \in [K] \backslash \mathcal{T})$. By definition, the probability of error in reconstructing all missing IVAs tends to 0 as $T \to \infty$ if

$$\varlimsup_{P \to \infty} \varlimsup_{A \to \infty} \frac{A \cdot N}{T \cdot \binom{K}{r} \cdot \log P} < \frac{\text{Sum-DoF(r)}}{\binom{K}{r}(K-r)}. \tag{55}$$

We thus conclude that

$$\Delta^\star(r) \leq \lim_{P \to \infty} \lim_{A \to \infty} \frac{T}{A \cdot K \cdot N} \cdot \log P = \frac{K-r}{K} \frac{1}{\text{Sum-DoF(r)}}, \tag{56}$$

which proves the desired achievability result.

APPENDIX B
ANALYSIS OF IA SCHEME

By the choice of T in (30), the signal and interference space at Rx 1 is represented by the $T \times T$-matrix:

$$\mathbf{\Lambda}_1 = \Big[\underbrace{\mathbf{D}_1}_{\text{signal space}}, \underbrace{[\mathbf{W}_\mathcal{R}]_{\mathcal{R} \in [[K] \backslash \{1\}]^r}}_{\text{interference space}} \Big], \tag{57}$$

where the signal subspace is given by a collection of $T \times \left((K-2) \cdot \binom{K-2}{r-1} \cdot \eta^\Gamma\right)$-matrices

$$\mathbf{D}_1 \triangleq \Big[\mathbf{H}_{1,k} \mathbf{U}_\mathcal{R} \Big]_{\substack{k \in [K-1] \backslash \{1\}, \cdot \\ \mathcal{R} \in [[K] \backslash \{1\}]^r : \\ k \in \mathcal{R}}} \tag{58}$$

The signal space at Rx $j \in [K] \backslash \{1\}$ is represented by the $T \times \tilde{T}$ matrix:

$$\mathbf{\Lambda}_j \triangleq \Big[\underbrace{\mathbf{D}_j}_{\text{signal space}}, \underbrace{[\mathbf{W}_\mathcal{R}]_{\mathcal{R} \in [[K] \backslash \{1\}]^r : \ j \notin \mathcal{R}}}_{\text{interference space}} \Big]. \tag{59}$$

where the signal subspace $\mathbf{D}_j$ is given by the collection of $T \times \left(r \cdot \binom{K-1}{r} \cdot \eta^\Gamma\right)$-matrices

$$\mathbf{D}_j \triangleq [\mathbf{H}_{j,k} \mathbf{U}_\mathcal{R}]_{\substack{\mathcal{R} \in [[K] \backslash \{1\}]^r : \\ j \in \mathcal{R}, \\ k \in [K] \backslash \{j\}}} \tag{60}$$

and

$$\tilde{T} \triangleq r \cdot \binom{K-1}{r} \cdot \eta^\Gamma + \binom{K-2}{r} \cdot (\eta+1)^\Gamma. \tag{61}$$

According to Lemmas 2 and 3 below, $\{\mathbf{\Lambda}_j\}$ is full column-rank if each column has different exponent vector $\boldsymbol{\alpha}$, which follows by the way we constructed the matrices $\mathbf{U}_\mathcal{R}$ and $\mathbf{W}_\mathcal{R}$. Indeed:

- For each $\mathcal{R} \in [[K] \backslash \{1\}]^r$, matrices $\mathbf{U}_\mathcal{R}$ and $\mathbf{W}_\mathcal{R}$ are constructed using a dedicated i.i.d. vector $\mathbf{\Xi}_\mathcal{R}$ that is independent of all other random variables in the system and thus the vectors $\mathbf{\Xi}_\mathcal{R}$ can play the roles of the vectors $\mathbf{\Xi}_i$ in Lemma 3.
- For each term $\mathbf{H}\mathbf{U}_\mathcal{R}$ in (58) and (60), we have $\mathbf{H} \notin \mathcal{H}_\mathcal{R}$. Thus $\mathbf{H}$ is not used in the construction of neither $\mathbf{U}_\mathcal{R}$ nor $\mathbf{W}_\mathcal{R}$ and induces a unique exponent on the corresponding columns in the signal space which is 0 in all columns of the interference space $\mathbf{W}_\mathcal{R}$.

This proves that based on the "cleaned" signal (40), each receiving node j can separate the various desired signals from each other as well as from the non-desired interfering signals. Since each codeword $\mathbf{b}_{k,\mathcal{T}}^j$ occupies η^Γ dimensions out of the T dimensions, we obtain that whenever

$$\frac{|\mathbf{b}_{k,\mathcal{T}}^j|}{T} \leq \frac{\eta^\Gamma}{T} \log P + o(\log P), \tag{62}$$

for an appropriate function $o(\log P)$ that grows slowlier than $\log P$, each codeword $\mathbf{b}_{k,\mathcal{T}}^j$ can be decoded with arbitrary small probability of error as $\eta \to \infty$.

Since $(K-2) \cdot \binom{K-2}{r-1}$ codewords are sent to Node 1, and $r\binom{K-1}{r}$ codewords to any other Node $j = 2, \ldots, K$, and since

$$\lim_{\eta \to \infty} \frac{\eta^\Gamma}{T} = \frac{1}{(K-2)\binom{K-2}{r-1} + \binom{K-1}{r}}, \tag{63}$$

we conclude that a sum-DoF of

$$\begin{aligned}
\text{Sum-DoF} &= \frac{(K-2) \cdot \binom{K-2}{r-1} + (K-1)r\binom{K-1}{r}}{(K-2)\binom{K-2}{r-1} + \binom{K-1}{r}} \\
&= \frac{r(K-1)^2 + r(K-2)}{r(K-2) + K - 1}
\end{aligned} \tag{64}$$

is achievable over the system. This establishes the desired achievability result.

Lemma 2. *Let $\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_m$ be independent random vectors with i.i.d. entries drawn according to continuous distributions. for any $L \leq m$ and L different exponent vectors*

$$\boldsymbol{\alpha}_j = (\alpha_{j,1}, \ldots, \alpha_{j,m}) \in \mathbb{Z}_+^m, \quad j \in [L],$$

the $m \times L$ matrix $\mathbf{M}$ with row-i and column-j entry

$$M_{i,j} = \prod_{k=1}^{m} (s_i)^{\alpha_{j,k}}, \quad i \in [m], \ j \in [L], \quad (65)$$

is full rank almost surely.

Lemma 3. *Consider numbers $\{n_1, n_2, \cdots, n_{\tilde{K}}\} \in \mathbb{Z}_+^{\tilde{K}}$ so that their sum $C \triangleq \sum_{i=1}^{\tilde{K}} n_i \leq \mathsf{T}$. Assume that for each $i \in [\tilde{K}]$ and $k \in [n_i]$, $\mathbf{B}_{i,k} \in \mathbb{C}^{\mathsf{T} \times \mathsf{T}}$ is a diagonal matrix so that all square sub-matrices of the following matrices $\{\mathbf{B}_i\}_{i \in \tilde{K}}$ are full rank:*

$$\mathbf{B}_i \triangleq \left[\mathbf{B}_{i,1} \cdot \mathbf{1}_\mathsf{T}, \mathbf{B}_{i,2} \cdot \mathbf{1}_\mathsf{T}, \cdots, \mathbf{B}_{i,n_i} \cdot \mathbf{1}_\mathsf{T} \right], \quad i \in [\tilde{K}], \quad (66)$$

where $\mathbf{1}_\mathsf{T}$ denotes a T-dimensional all-one column vector.

Let further $\{\boldsymbol{\Xi}_i\}_{i \in \tilde{K}}$ be independent T-vectors with entries drawn i.i.d. from continuous distributions and define the $\mathsf{T} \times n_i$-matrices

$$\mathbf{A}_i \triangleq \left[\mathbf{B}_{i,1} \cdot \boldsymbol{\Xi}_i, \mathbf{B}_{i,2} \cdot \boldsymbol{\Xi}_i, \cdots, \mathbf{B}_{i,n_i} \cdot \boldsymbol{\Xi}_i \right], \quad i \in [\tilde{K}]. \quad (67)$$

Then, the $\mathsf{T} \times C$-matrix

$$\boldsymbol{\Lambda} \triangleq \left[\mathbf{A}_1, \mathbf{A}_2, \cdots, \mathbf{A}_{\tilde{K}} \right] \quad (68)$$

has full column rank almost surely.

Proof: We assume that the matrix $\boldsymbol{\Lambda}$ is a square matrix i.e. $C = \mathsf{T}$. If $\mathsf{T} > C$, we take a square submatrix of $\boldsymbol{\Lambda}$ and perform the same proof steps on the submatrix.

Define

$$F \left(\boldsymbol{\Xi}_1, \ldots, \boldsymbol{\Xi}_{\tilde{K}} \right) \triangleq \det(\boldsymbol{\Lambda}) \quad (69)$$

which is a polynomial of $\boldsymbol{\Xi}_1, \boldsymbol{\Xi}_2, \cdots, \boldsymbol{\Xi}_{\tilde{K}}$ as the determinant is a polynomial of the entries of $\boldsymbol{\Lambda}$.

For the vectors

$$\boldsymbol{d}_i = [\ \underbrace{0, \cdots 0,}_{(n_1 + \cdots + n_{i-1}) \ \mathrm{0s}} \ \underbrace{1, \cdots 1,}_{n_i \ \mathrm{1s}} \ \underbrace{0, \cdots 0}_{(n_{i+1} + \cdots + n_{\tilde{K}}) \ \mathrm{0s}} \ ]^T, \quad i \in \tilde{K}, \quad (70)$$

the polynomial evaluates to

$$F \left(\boldsymbol{d}_1, \ldots, \boldsymbol{d}_{\tilde{K}} \right) = \det \begin{pmatrix} \mathbf{B}'_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{B}'_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{B}'_{\tilde{K}} \end{pmatrix} \quad (71)$$

$$= \prod_{i=1}^{\tilde{K}} \det(\mathbf{B}'_i) \neq 0 \quad (72)$$

where $\mathbf{B}'_i$ is the $n_i \times n_i$ square sub-matrix of $\mathbf{B}_i$ consisting of its rows $(n_1 + \cdots + n_{i-1} + 1)$ to $(n_1 + \cdots + n_{i-1} + n_i)$. The inequality holds by our assumption that all square sub-matrices of $\mathbf{B}_i$ are full rank.

We conclude that F is a non-zero polynomial and thus $F \left(\boldsymbol{\Xi}_1, \ldots, \boldsymbol{\Xi}_{\tilde{K}} \right)$ equals 0 with probability 0 because the entries of $\boldsymbol{\Xi}_1, \boldsymbol{\Xi}_2, \cdots, \boldsymbol{\Xi}_{\tilde{K}}$ are drawn independently from continuous distributions. ∎

APPENDIX C
PROOF OF THE NDT LOWER BOUND IN THEOREM 1

Consider a fixed file assignment (map phase), and for any positive power P a sequence (in T) of wireless distributed computing systems satisfying (9) for the given file assignment. (Since for finite N there are only a finite number of different file assignments irrespective of P and T, we can fix the assignment.) The following limiting behaviour must hold.

Lemma 4. *Consider two disjoint sets $\mathcal{T}$ and $\mathcal{R}$ of same size*

$$|\mathcal{T}| = |\mathcal{R}|, \quad (73)$$

and define $\mathcal{F} \triangleq [\mathsf{K}] \backslash (\mathcal{R} \cup \mathcal{T})$. Let $\mathcal{M} \subseteq [\mathsf{N}]$ be the set of files known only to nodes $\mathcal{T}$ but not to any other node and partition the set of all IVAs $\mathcal{A}$ it into the following disjoint subsets:

$$\mathcal{W}_r \triangleq \{a_{j,m}\}_{\substack{j \in \mathcal{R} \\ m \in [\mathsf{N}] \backslash \mathcal{M}_j}}, \quad (74)$$

$$\mathcal{W}_t \triangleq \{a_{j,m}\}_{\substack{j \in (\mathcal{T} \cup \mathcal{F}) \\ m \in \mathcal{M} \backslash \mathcal{M}_j}}. \quad (75)$$

For any sequence of distributed computing systems:

$$d \triangleq \varlimsup_{P \to \infty} \varlimsup_{\mathsf{T} \to \infty} \frac{\mathsf{A}}{\mathsf{T} \log \mathsf{P}} \leq \frac{|\mathcal{T}|}{|\mathcal{W}_t| + |\mathcal{W}_r|} \quad (76)$$

(Notice that $\mathcal{W}_r$ denotes the set of all IVAs intended to nodes in $\mathcal{R}$ and $\mathcal{W}_t$ the set of IVAs deduced from files in $\mathcal{M}$ and intended for nodes not in $\mathcal{R}$.)

Proof. Denote by $\mathcal{H}$ the set of all channel coefficients to all nodes in the system and define $\mathcal{W}_c \triangleq \mathcal{A} \backslash (\mathcal{W}_r \cup \mathcal{W}_t)$. Since channel coefficients and IVAs are independent, we have

$$H(\mathcal{W}_t, \mathcal{W}_r)$$
$$= H(\mathcal{W}_t, \mathcal{W}_r | \mathcal{W}_c, \mathcal{H}) \quad (77)$$
$$= I(\mathcal{W}_t, \mathcal{W}_r; \mathbf{Y}_\mathcal{R} | \mathcal{W}_c, \mathcal{H}) + H(\mathcal{W}_t, \mathcal{W}_r | \mathcal{W}_c, \mathbf{Y}_\mathcal{R}, \mathcal{H}) \quad (78)$$
$$= h(\mathbf{Y}_\mathcal{R} | \mathcal{W}_c, \mathcal{H}) - h(\mathbf{Z}_\mathcal{R})$$
$$\quad + H(\mathcal{W}_r | \mathcal{W}_c, \mathbf{Y}_\mathcal{R}, \mathcal{H}) + H(\mathcal{W}_t | \mathcal{W}_r, \mathcal{W}_c, \mathbf{Y}_\mathcal{R}, \mathcal{H}) \quad (79)$$
$$\leq h(\mathbf{Y}_\mathcal{R} | \mathcal{W}_c, \mathcal{H}) - h(\mathbf{Z}_\mathcal{R})$$
$$\quad + \mathsf{T}\epsilon_\mathsf{T} + H(\mathcal{W}_t | \mathcal{W}_r, \mathcal{W}_c, \mathbf{Y}_\mathcal{R}, \mathcal{H}), \quad (80)$$

where we defined $\mathbf{Y}_\mathcal{A} \triangleq [\mathbf{Y}_j]_{j \in \mathcal{A}}$ for a set $\mathcal{A} \subseteq [\mathsf{K}]$ and ϵ_T is a vanishing sequence as $\mathsf{T} \to \infty$. Here the inequality holds by Fano's inequality, because $\mathcal{W}_r$ is decoded from $\mathbf{Y}_\mathcal{R}$ and $\mathcal{W}_c$, and because we impose vanishing probability of error (9).

Again by Fano's inequality and by (9), there exists a vanishing sequence ϵ'_T such that

$$H(\mathcal{W}_t | \mathcal{W}_r, \mathcal{W}_c, \mathbf{Y}_\mathcal{R}, \mathcal{H})$$
$$\leq I(\mathcal{W}_t; \mathbf{Y}_{(\mathcal{F} \cup \mathcal{T})} | \mathcal{W}_r, \mathcal{W}_c, \mathbf{Y}_\mathcal{R}, \mathcal{H}) + \mathsf{T}\epsilon'_\mathsf{T} \quad (81)$$
$$= h(\mathbf{Y}_{(\mathcal{F} \cup \mathcal{T})} | \mathcal{W}_r, \mathcal{W}_c, \mathbf{Y}_\mathcal{R}, \mathcal{H}) \quad (82)$$
$$\quad - h(\mathbf{Y}_{(\mathcal{F} \cup \mathcal{T})} | \mathcal{W}_r, \mathcal{W}_t, \mathcal{W}_c, \mathbf{Y}_\mathcal{R}, \mathcal{H}) + \mathsf{T}\epsilon'_\mathsf{T} \quad (83)$$
$$\leq h(\bar{\mathbf{Y}}_{(\mathcal{F} \cup \mathcal{T})} | \bar{\mathbf{Y}}_\mathcal{R}, \mathcal{H}) - h(\mathbf{Z}_{(\mathcal{F} \cup \mathcal{T})}) + \mathsf{T}\epsilon'_\mathsf{T}, \quad (84)$$

where $\bar{\mathbf{Y}}_\mathcal{A} \triangleq [\bar{\mathbf{Y}}_j]_{j \in \mathcal{A}}$ and $\bar{\mathbf{Y}}_j$ denotes Node j's "cleaned" signal without the inputs that do not depend on files in $\mathcal{M}$ but only on IVAs $\mathcal{W}_r \cup \mathcal{W}_c$:

$$\bar{\mathbf{Y}}_j \triangleq \mathbf{H}_{j,\mathcal{T}} \mathbf{X}_\mathcal{T} + \mathbf{Z}_j, \quad j \in \mathcal{T} \cup \mathcal{F}.$$

Here, $\mathbf{H}_{\mathcal{A},\mathcal{B}}$ denotes the channel matrix from set $\mathcal{B}$ to set $\mathcal{A}$.

To bound the first term in (84), we introduce a random variable E indicating whether the matrix $\mathbf{H}_{\mathcal{R},\mathcal{T}}$ is invertible ($E = 1$) or not ($E = 0$). If this matrix is invertible and $E = 1$, then the input vector $\mathbf{X}_{\mathcal{T}}$ can be computed from $\bar{\mathbf{Y}}_{\mathcal{R}}$ up to noise terms. Based on this observation and defining the residual noise terms

$$\bar{\mathbf{Z}}_j \triangleq \mathbf{Z}_j - \mathbf{H}_{j,\mathcal{T}}\mathbf{H}_{\mathcal{R},\mathcal{T}}^{-1}\mathbf{Z}_{\mathcal{R}}, \quad \text{if } E = 1, \qquad (85)$$

we obtain:

$$h(\bar{\mathbf{Y}}_{(\mathcal{F}\cup\mathcal{T})}|\bar{\mathbf{Y}}_{\mathcal{R}}, \mathcal{H})$$
$$\leq \mathbb{P}(E = 1) \cdot h\left(\bar{\mathbf{Z}}_{(\mathcal{F}\cup\mathcal{T})}|\bar{\mathbf{Y}}_{\mathcal{R}}, \mathcal{H}, E = 1\right)$$
$$+ \mathbb{P}(E = 0) \cdot h\left(\bar{\mathbf{Y}}_{(\mathcal{F}\cup\mathcal{T})}|\bar{\mathbf{Y}}_{\mathcal{R}}, \mathcal{H}, E = 0\right) \qquad (86)$$
$$\leq h\left(\bar{\mathbf{Z}}_{(\mathcal{F}\cup\mathcal{T})}\right) + \mathbb{P}(E = 0) h\left(\bar{\mathbf{Y}}_{(\mathcal{F}\cup\mathcal{T})}|\bar{\mathbf{Y}}_{\mathcal{R}}, \mathcal{H}, E = 0\right). \quad (87)$$

Since the channel coefficients follow continuous distribution, $\mathbf{H}_{\mathcal{R},\mathcal{T}}$ is invertible almost surely, implying $\mathbb{P}(E = 0) = 0$. By the boundedness of the entropy term $h\left(\bar{\mathbf{Y}}_{(\mathcal{F}\cup\mathcal{T})}|\bar{\mathbf{Y}}_{\mathcal{R}}, \mathcal{H}, E = 0\right)$ (since power P and channel coefficients are bounded), this implies

$$h(\bar{\mathbf{Y}}_{(\mathcal{F}\cup\mathcal{T})}|\bar{\mathbf{Y}}_{\mathcal{R}}, \mathcal{H}) \leq h(\bar{\mathbf{Z}}_{(\mathcal{F}\cup\mathcal{T})}),$$

which combined with (80) and (84) yields:

$$H(\mathcal{W}_t, \mathcal{W}_r) \leq h(\mathbf{Y}_{\mathcal{R}}|\mathcal{H}) - h(\mathbf{Z}_{\mathcal{R}}) + h(\bar{\mathbf{Z}}_{(\mathcal{F}\cup\mathcal{T})})$$
$$- h(\mathbf{Z}_{(\mathcal{F}\cup\mathcal{T})}) + \mathsf{T}(\epsilon_{\mathsf{T}} + \epsilon'_{\mathsf{T}})$$
$$\leq \mathsf{T}|\mathcal{R}|\log(\mathsf{P}) + \mathsf{T}C_{\mathsf{T},\mathcal{H}}, \qquad (88)$$

where $C_{\mathsf{T},\mathcal{H}}$ is a function that is uniformly bounded over all realizations of channel matrices and powers P. Noticing

$$H(\mathcal{W}_t, \mathcal{W}_r) = \mathsf{A}(|\mathcal{W}_t| + |\mathcal{W}_r|), \qquad (89)$$

dividing (88) by $\mathsf{T}\log(\mathsf{P})$, and letting $\mathsf{P} \to \infty$, establishes the lemma because $|\mathcal{R}| = |\mathcal{T}|$ and $\mathsf{T}C_{\mathsf{T},\mathcal{H}}$ is bounded. $\qquad \square$

For each subset $\mathcal{T} \subseteq [\mathsf{K}]$, let $\mathcal{B}_{\mathcal{T}}^j$ denote the set of IVAs that are computed exclusively at nodes in set $\mathcal{T}$ and intended for reduce function j. Define $b_{\mathcal{T}} = |\mathcal{B}_{\mathcal{T}}^j|$, which does not depend on the index of the reduce function $j \in [\mathsf{K}]\backslash\mathcal{T}$.

Choose two disjoint subsets $\mathcal{T}$ and $\mathcal{R}$ of same size $|\mathcal{T}| = |\mathcal{R}|$. By Lemma 4, and rewriting the sets $\mathcal{W}_t$ and $\mathcal{W}_r$ in the lemma in terms of the sets $\{\mathcal{B}_{\mathcal{T}}^j\}$, we obtain:

$$\frac{|\mathcal{T}|}{d} \geq \sum_{\mathcal{T}\subseteq[\mathsf{K}]} \sum_{j\in\mathcal{R}\backslash\mathcal{T}} |\mathcal{B}_{\mathcal{T}}^j| + \sum_{\mathcal{G}\subseteq\mathcal{T}} \sum_{j\in[\mathsf{K}]\backslash(\mathcal{R}\cup\mathcal{G})} |\mathcal{B}_{\mathcal{G}}^j| \quad (90)$$

$$= \sum_{\mathcal{T}\subseteq[\mathsf{K}]} |\mathcal{R}\backslash\mathcal{T}| \cdot b_{\mathcal{T}} + \sum_{\mathcal{G}\subseteq\mathcal{T}} (\mathsf{K} - |\mathcal{R}| - |\mathcal{G}|) \cdot b_{\mathcal{G}}. \quad (91)$$

Summing up Equality (91) over all sets $\mathcal{T}$ and $\mathcal{R}$ of constant size $t \leq \mathsf{K}/2$, we obtain:

$$\binom{\mathsf{K}}{t} \cdot \binom{\mathsf{K}-t}{t} \cdot \frac{t}{d}$$
$$\geq \sum_{\mathcal{T}\in[[\mathsf{K}]]^t} \sum_{\mathcal{R}\in[[\mathsf{K}]\backslash\mathcal{T}]^t} \sum_{\mathcal{T}\subseteq[\mathsf{K}]} |\mathcal{R}\backslash\mathcal{T}| \cdot b_{\mathcal{T}}$$

$$+ \sum_{\mathcal{T}\in[[\mathsf{K}]]^t} \sum_{\mathcal{R}\in[[\mathsf{K}]\backslash\mathcal{T}]^t} \sum_{\mathcal{G}\subseteq\mathcal{T}} (\mathsf{K} - t - |\mathcal{G}|) \cdot b_{\mathcal{G}} \qquad (92)$$

$$= \sum_{\mathcal{T}\subseteq[\mathsf{K}]} \binom{\mathsf{K}}{t} \cdot \binom{\mathsf{K}-t}{t} \cdot (\mathsf{K} - |\mathcal{T}|)\frac{t}{\mathsf{K}} b_{\mathcal{T}}$$

$$+ \sum_{\substack{\mathcal{G}\subseteq[\mathsf{K}]: \\ |\mathcal{G}|\leq t}} \binom{\mathsf{K}-|\mathcal{G}|}{t-|\mathcal{G}|} \cdot \binom{\mathsf{K}-t}{t} \cdot (\mathsf{K} - t - |\mathcal{G}|)b_{\mathcal{G}} \quad (93)$$

$$= \binom{\mathsf{K}}{t} \cdot \binom{\mathsf{K}-t}{t} \cdot t\left(\mathsf{N} - \frac{\mathsf{r}\mathsf{N}}{\mathsf{K}}\right)$$

$$+ \sum_{\substack{\mathcal{G}\subseteq[\mathsf{K}]: \\ |\mathcal{G}|\leq t}} \binom{\mathsf{K}-|\mathcal{G}|}{t-|\mathcal{G}|} \cdot \binom{\mathsf{K}-t}{t} \cdot (\mathsf{K} - t - |\mathcal{G}|)b_{\mathcal{G}}, \quad (94)$$

where we define $\binom{a}{0} = 1$ for any positive integer a. The first equality holds because for a given set $\mathcal{T}$, each element of $[\mathsf{K}]\backslash\mathcal{T}$ is present in a fraction of t/K pairs of the admissible sets $(\mathcal{R}, \mathcal{T})$ and the last equality holds because

$$\sum_{\mathcal{T}\subseteq[\mathsf{K}]} b_{\mathcal{T}} = \mathsf{N}, \quad \sum_{\mathcal{T}\subseteq[\mathsf{K}]} |\mathcal{T}| \cdot b_{\mathcal{T}} \leq \mathsf{r} \cdot \mathsf{N}. \qquad (95)$$

Dividing both sides of (94) by $\binom{\mathsf{K}}{t}\binom{\mathsf{K}-t}{t}t$, and defining $b_i \triangleq \sum_{\mathcal{T}\in[[\mathsf{K}]]^i} b_{\mathcal{T}}$, for any $t \in [\lfloor\mathsf{K}/2\rfloor]$ we obtain:

$$\frac{1}{d} \geq \mathsf{N} - \frac{\mathsf{r}\cdot\mathsf{N}}{\mathsf{K}} + \min_{\substack{b_1,\ldots,b_{\mathsf{K}}\in\mathbb{Z}^+: \\ \sum_{i=1}^{\mathsf{K}} b_i=\mathsf{N} \\ \sum_{i=1}^{\mathsf{K}} ib_i\leq\mathsf{r}\mathsf{N}}} \sum_{i=1}^{t} C_t(i)b_i, \quad t \in [\lfloor\mathsf{K}/2\rfloor], \quad (96)$$

where $C_t(i)$ is defined in (46).

For any $t \in [\lfloor\mathsf{K}/2\rfloor]$, the sequence of coefficients $C_t(1), C_t(2), \ldots, C_t(t)$ is convex and non-increasing, see Appendix D. Based on this convexity, it can be shown (see Appendix E) that for any $\mathsf{r} < t + 1$ there exists a solution to the minimization problem in (96) putting only positive masses on $b_{\lfloor\mathsf{r}\rfloor}^*$ and $b_{\lceil\mathsf{r}\rceil}^*$ in the unique way satisfying

$$b_{\lfloor\mathsf{r}\rfloor}^* + b_{\lceil\mathsf{r}\rceil}^* = \mathsf{N} \qquad (97)$$
$$\lfloor\mathsf{r}\rfloor b_{\lfloor\mathsf{r}\rfloor}^* + \lceil\mathsf{r}\rceil b_{\lceil\mathsf{r}\rceil}^* = \mathsf{r}\mathsf{N}. \qquad (98)$$

For $\mathsf{r} \geq t + 1$ an optimal solution consists of setting $b_{\lfloor\mathsf{r}\rfloor}^* = \mathsf{N}$, in which case the minimization in (96) evaluates to 0.

For $\mathsf{r} \geq 2$, the lower bound on the NDT in the theorem is then obtained by plugging these optimum values into bound (96) for the choice $t = \lfloor\mathsf{K}/2\rfloor$. For $\mathsf{r} = 1$ we choose $t = 1$, and for $\mathsf{r} \in (1, 2)$ we maximize over the value of t.

APPENDIX D
PROOF OF MONOTONICITY AND CONVEXITY OF VALUES $C_i^{(t)}$

We shall prove monotonicity and convexity of the values

$$D_i^{(t)} \triangleq \frac{\mathsf{K}!}{t!}t \cdot C_i^{(t)} \qquad (99)$$
$$= \frac{(\mathsf{K}-i)!}{(t-i)!}(\mathsf{K}-t-i), \qquad i \in [t]. \qquad (100)$$

Notice that

$$D_{i-1}^{(t)} = D_i^{(t)} \frac{\mathsf{K} - i + 1}{t - i + 1} \cdot \frac{\mathsf{K} - t - i + 1}{\mathsf{K} - t - i} \qquad (101)$$

$$D_{i+1}^{(t)} = D_i^{(t)} \frac{t - i}{\mathsf{K} - i} \cdot \frac{\mathsf{K} - t - i - 1}{\mathsf{K} - t - i}, \qquad (102)$$

and the monotonicity

$$D_{i-1}^{(t)} > D_i^{(t)} \qquad (103)$$

simply follows because $K > t$ implies

$$\mathsf{K} - i + 1 > t - i + 1 \quad \text{and} \quad \mathsf{K} - t - i + 1 > \mathsf{K} - t - i. \qquad (104)$$

To prove convexity, we shall prove that

$$D_{i+1}^{(t)} + D_{i-1}^{(t)} \geq 2 D_i^{(t)}, \qquad (105)$$

or equivalently

$$\frac{t - i}{\mathsf{K} - i} \cdot \frac{\mathsf{K} - t - i - 1}{\mathsf{K} - t - i} + \frac{\mathsf{K} - i + 1}{t - i + 1} \cdot \frac{\mathsf{K} - t - i + 1}{\mathsf{K} - t - i} \geq 2. \quad (106)$$

Multiplying both sides with the denominators, we see that this condition is equivalent to

$$\begin{aligned}(t - i)(\mathsf{K} - t - i - 1)(t - i + 1) \\ + (\mathsf{K} - i + 1)(\mathsf{K} - t - i + 1)(\mathsf{K} - i) \\ \geq 2(\mathsf{K} - i)(t - i + 1)(\mathsf{K} - t - i). \qquad (107)\end{aligned}$$

and after rearranging the terms:

$$\begin{aligned}(\mathsf{K} - t)(\mathsf{K} - i)(\mathsf{K} - t - i) + (\mathsf{K} - i)(\mathsf{K} - i + 1) \\ \geq (\mathsf{K} - t)(t - i + 1)(\mathsf{K} - t - i) + (t - i)(t - i + 1). \quad (108)\end{aligned}$$

This last inequality is easily verified by noting that $\mathsf{K} \geq t + 1$.

APPENDIX E
PROOF OF STRUCTURE OF MINIMIZER

Start with any feasible vector $b_1, \ldots, b_\mathsf{K}$ and consider two indices $i < j$ with non-zero masses, $b_i > 0$ and $b_j > 0$. Updating this vector as

$$b_i' = b_i - \Delta, \qquad \text{and} \qquad b_{i+1}' = b_{i+1} + \Delta, \quad (109)$$

$$b_{j-1}' = b_{j-1} + \Delta, \qquad \text{and} \qquad b_j' = b_j - \Delta, \quad (110)$$

for any $\Delta \in [0, \min\{b_i, b_j\}]$, results again in a feasible solution vector, which has smaller objective function due to the convexity of the coefficients $\{C_i^{(t)}\}$.

Applying this argument iteratively, one can conclude that there must exist an optimal solution vector where all entries are zero except for two masses $b_k > 0$ and $b_{k+1} \geq 0$. Since $\sum_{i=1}^{\mathsf{K}} i b_i \leq \mathsf{r} \mathsf{N}$, the index k cannot exceed r. By the decreasing monotonicity of the coefficients $C_i^{(t)}$, the optimal solution must then be to choose $b_{\lfloor \mathsf{r} \rfloor} > 0$ and $b_{\lfloor \mathsf{r} \rfloor + 1} \geq 0$ and all other masses equal to 0. Since there is a unique such choice satisfying $\sum_{i=1}^{\mathsf{K}} i b_i \leq \mathsf{r} \mathsf{N}$ and $\sum_{i=1}^{\mathsf{K}} b_i = \mathsf{N}$, this concludes the proof.