

Necessity of Cooperative Transmissions for Wireless MapReduce

Yue Bi*, Michèle Wigger*

* *LTCI, Telecom Paris, IP Paris, 91120 Palaiseau, France; biyue111@gmail.com, michele.wigger@telecom-paris.fr*

Abstract—The paper presents an improved upper bound (achievability result) on the optimal tradeoff between Normalized Delivery Time (NDT) and computation load for distributed computing MapReduce systems in certain ranges of the parameters. The upper bound is based on interference alignment combined with zero-forcing. The paper further provides a lower bound (converse) on the optimal NDT-computation tradeoff that can be achieved when intermediate values (IVA) are partitioned into sub-IVAs, and these sub-IVAs are then transmitted (in an arbitrary form) by a single node, without cooperation among nodes. For appropriate linear functions (e.g., XORs), such non-cooperative schemes can achieve some of the best NDT-computation tradeoff points so far obtained in the literature. However, as our lower bound shows, any non-cooperative scheme achieves a worse NDT-computation tradeoff than our new proposed scheme for certain parameters, thus proving the necessity of cooperative schemes like zero-forcing to attain the optimal NDT-computation tradeoff.

Index Terms—Wireless distributed computing, interference alignment, zero-forcing, NDT.

I. INTRODUCTION

MapReduce systems are becoming important building blocks also in wireless scenarios because they allow us to distribute heavy computation tasks over multiple distributed nodes. Important performance parameters of such systems are the *computation load* r , which characterizes the average amount of data stored at the nodes (i.e., memory requirements), and the *Normalized Delivery Time (NDT)* Δ , which characterizes the duration of the wireless shuffle phase (i.e., communication requirements). The optimal tradeoff between memory requirements and communication duration was first characterized in [1], but under the assumption of perfect multicast communication links. The optimal scheme for this noiseless setup splits intermediate values (IVA) into message parts and then communicates appropriate XORs of these parts. Decoding requires that receiving nodes use their knowledge of the message parts transmitted by other nodes to cancel non-desired parts of the XORs. Similar assumptions of ideal multicast or noiseless communication links have also been adopted in several related works, e.g., [2]–[4].

Several subsequent works [5]–[8] have studied the *NDT-computation tradeoff* of wireless MapReduce, i.e., the smallest NDT that can be achieved under a constraint on the computation load. The problem was first investigated for wireless interference networks in [5], which proposed a coding scheme based on one-shot beamforming and zero-forcing (ZF). In our previous works in [6], [7], we obtained improved NDT-computation tradeoffs by introducing interference alignment

(IA) [9], [10] to the shuffle phase. For computation loads $r = \lfloor (K-1)/2 \rfloor$, where K denotes the number of nodes and is assumed to be odd, our work in [7] also proposes an improved IA scheme based on ZF. Notice that ZF introduces complex dependencies between IA matrices, rendering performance analysis of IA schemes involved [11].

In this paper, we improve our IA & ZF scheme in [7]. Our new scheme applies to an arbitrary number of nodes K but still requires that $r = \lfloor (K-1)/2 \rfloor$. For odd values of K it strictly improves over our previous scheme in [7]. The new scheme is again based on IA & ZF and the analysis requires proving algebraic independence of certain functions. We provide an analytic proof for $K = 5$ in this paper as well as numeric verification for $K \in \{6, \dots, 15\}$.

Inspired by the excellent performance of the scheme in [1] and to avoid complicated analyses, several schemes (also for related cache-aided scenarios) [6], [7], [12] focused on exploiting the multicast behaviour of wireless MapReduce communication and the IVA side-information at the nodes, but without any form of cooperation between transmit signals. Specifically, the schemes divide the IVAs into submessages and each submessage is transmitted (in an arbitrary form) by only one of the nodes, without cooperation from the other nodes. Transmit signals across nodes are thus independent of each other. Nodes still use all their side-information about IVAs in their decoding steps. In this work, we provide a lower bound (converse result) on the NDT-computation tradeoff that can be achieved by any such non-cooperative scheme. We further show that this lower bound lies above the NDT-computation tradeoffs obtained by our new IA & ZF scheme, thus proving strict suboptimality of non-cooperative schemes. Notice that the previously obtained NDT-computation tradeoffs do not improve over the non-cooperative lower bound and do not allow us to obtain the desired conclusion.

Notations: We use standard notation, and also define $[n] \triangleq \{1, 2, \dots, n\}$ and $[\mathcal{A}]^t$ as the collection of all the subsets of \mathcal{A} with cardinality t . i.e. $[\mathcal{A}]^t \triangleq \{\mathcal{T} : \mathcal{T} \subseteq \mathcal{A}, |\mathcal{T}| = t\}$.

II. WIRELESS MAPREDUCE FRAMEWORK

Consider a wireless MapReduce system with $K \geq 5$ nodes labeled $1, \dots, K$, with M input files W_1, \dots, W_M , and with K output functions ϕ_1, \dots, ϕ_K , where function ϕ_k is assigned to Node k . The system follows the MapReduce framework, in which each output function is decomposed as

$$\phi_q(W_1, \dots, W_M) = v_q(a_{q,1}, \dots, a_{q,M}), \quad q \in [K], \quad (1)$$

where

$$a_{q,p} = u_{q,p}(W_p), \quad p \in [M], \quad (2)$$

denotes the intermediate value (IVA) computed from file W_p . All IVAs are assumed independent and consist of A i.i.d. bits.

The MapReduce framework has 3 phases:

a) *Map phase*: Each node k is assigned a subset of files $\mathcal{M}_k \subseteq [M]$ and computes the IVAs $\{a_{q,p} : p \in \mathcal{M}_k, q \in [K]\}$.

b) *Shuffle phase*: Nodes exchange IVAs over n uses of a wireless channel in full-duplex mode. Node k transmits

$$\mathbf{X}_k = f_k^{(n)}(\{a_{1,p}, \dots, a_{K,p}\}_{p \in \mathcal{M}_k}), \quad (3)$$

subject to the average power constraint

$$\frac{1}{n} \sum_{t=1}^n \mathbb{E}[|X_k(t)|^2] \leq P. \quad (4)$$

Node k observes

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

and $\mathbf{Y}_k \triangleq (Y_k(1), \dots, Y_k(n))$, where $\{H_{k,k'}(t)\}$ and $\{Z_k(t)\}$ are i.i.d. over time and across node pairs. Each coefficient $H_{p,q}(t)$ has independent and identically distributed real and imaginary parts, following a specified continuous distribution over a bounded interval $[-H_{\max}, H_{\max}]$ for some positive integer H_{\max} . All nodes have perfect channel state information.

Based on \mathbf{Y}_k and the IVAs computed locally, Node k decodes the missing IVAs as

$$\hat{a}_{k,p} = g_{k,p}^{(n)}(\{a_{q,i}\}_{i \in \mathcal{M}_k, q \in [K]}, \mathbf{Y}_k), \quad p \notin \mathcal{M}_k. \quad (6)$$

c) *Reduce phase*: Each node $k \in [K]$ applies its reduce function v_k to its available IVAs to compute its assigned output function ϕ_k .

The performance of the distributed computing system is measured by two factors: *computation load* and *normalized delivery time (NDT)*. The computation load is defined as

$$r \triangleq \sum_{k \in [K]} \frac{|\mathcal{M}_k|}{M}, \quad (7)$$

and NDT is defined as

$$\Delta \triangleq \lim_{P \rightarrow \infty} \lim_{A \rightarrow \infty} \frac{n}{AKM} \log P. \quad (8)$$

We define the optimal NDT-computation tradeoff $\Delta^*(r)$ as the infimum of all Δ for which there exist file assignments $\{\mathcal{M}_k\}$ and encoding/decoding functions satisfying

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

In this work, we choose a the well known combinatorial file assignment for the wireless MapReduce system. Specifically, the input files $\{W_1, \dots, W_M\}$ are partitioned into $\binom{K}{r}$ disjoint bundles, each associated with a size- r subset $\mathcal{T} \in \binom{[K]}{r}$ and stored at the nodes in \mathcal{T} . Since each file is stored at exactly r nodes, this assignment satisfies the computation load constraint.

A. Relation with r -fold cooperative channel

Under this combinatorial file assignment, the Shuffle phase naturally induces a wireless channel with r -fold transmitter cooperation. In particular, consider the wireless interference channel described in (5) under the average power constraint (4). For a given $r \in [K]$, any subset of r nodes $\mathcal{T} \in \binom{[K]}{r}$ jointly holds a message $a_{k,\mathcal{T}}$ intended for Node k , for each $k \in [K] \setminus \mathcal{T}$. Each message $a_{k,\mathcal{T}}$ is uniformly distributed over $\{1, \dots, 2^{nR_{k,\mathcal{T}}}\}$.

A rate tuple $(R_{k,\mathcal{T}} : \mathcal{T} \in \binom{[K]}{r}, k \in [K] \setminus \mathcal{T})$ is said to be achievable if there exist encoding and decoding functions such that the decoding error probabilities vanish as $n \rightarrow \infty$. The corresponding sum degrees of freedom is defined as

$$\text{SDoF}(r) \triangleq \lim_{P \rightarrow \infty} \sup_{\mathbf{R} \in \mathcal{C}(P)} \sum_{\mathcal{T} \in \binom{[K]}{r}} \sum_{k \in [K] \setminus \mathcal{T}} \frac{R_{k,\mathcal{T}}}{\log P}. \quad (10)$$

where the supremum is over all achievable rate tuples under power constraint P .

We have the following lemma, which establishes a direct connection between the NDT and the SDoF [7].

Lemma 1. *Under the combinatorial file assignment, for any $r \in [K]$,*

$$\Delta(r) = \left(1 - \frac{r}{K}\right) \frac{1}{\text{SDoF}(r)}. \quad (11)$$

III. MAIN RESULTS

Our first result is an achievable NDT-computation tradeoff that improves over previous results for $r = \lfloor (K-1)/2 \rfloor$.

Theorem 1 (Achievable NDT-computation tradeoff). *The NDT-computation tradeoff for $5 \leq K \leq 15^1$ and $r = \lfloor (K-1)/2 \rfloor$ satisfies*

$$\Delta(r) \leq \begin{cases} \frac{1}{K} \left(1 - \frac{r}{K}\right) \left(1 + \frac{4}{(K-2)^2}\right), & K \text{ even,} \\ \frac{1}{K} \left(1 - \frac{r}{K}\right) \left(1 + \frac{1}{(K-1)(K-2)}\right), & K \text{ odd.} \end{cases} \quad (12)$$

Proof: The coding scheme and corresponding SDoF for the r -fold cooperative channel is described in Section IV. Above tradeoff can then be obtained by applying Lemma 1. ■

As illustrated in Fig. 1, for $r = \lfloor (K-1)/2 \rfloor$, the proposed scheme strictly outperforms the schemes in [5]–[7] for all considered values of K .

Above achievability result is obtained using a coding scheme based on IA and ZF. A key-characteristics of ZF is that nodes do not send independent signals, but correlate them so as to avoid that certain signals interfere at some of the nodes. Our second result implies that under the proposed combinatorial file assignment, such cooperative transmission is essential to achieve the best NDT-computation tradeoff. In

¹The restriction $K \leq 15$ is due to the use of numerical verification. This choice does not suggest any inherent limitation of the proposed scheme for larger K , instead, The performance of the scheme has only been numerically validated for $K \leq 15$.

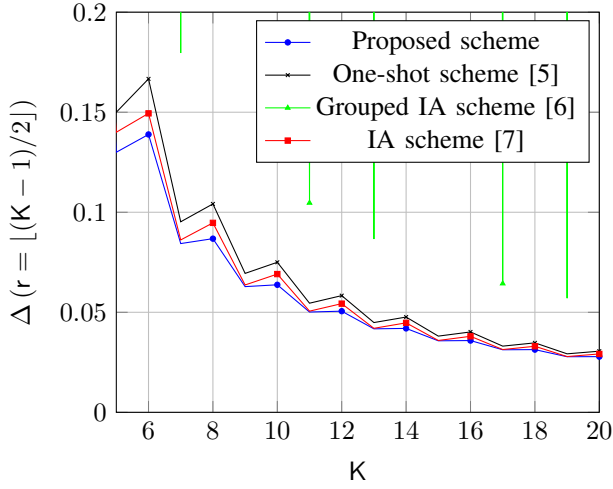


Fig. 1. Comparison of the achieved NDT-computation tradeoff between the proposed scheme and the schemes in [5]-[7].

the following Theorem 2, we provide a lower bound on the best NDT-computation tradeoff that is achievable using any kind of non-cooperative transmission. As we show, this lower bound lies above the NDT-computation tradeoff achieved by our IA and ZF scheme in Theorem 1.

Specifically, consider any MapReduce system with combinatorial file assignment and where each IVA $a_{q,p}$ is split into r equal length sub-messages $(a_{t,\mathcal{T}}^q)_{t \in \mathcal{T}}$ where \mathcal{T} is the set of nodes that have stored $a_{q,p}$, i.e., $\mathcal{T} = \{k: p \in \mathcal{M}_k\}$. Each node $t \in \mathcal{T}$ then transmits only $a_{t,\mathcal{T}}^q$ but not the other submessages implying node t 's transmit signal X_t^n is independent of the sub-messages $(a_{t',\mathcal{T}}^q)_{t' \in \mathcal{T} \setminus \{t\}}$. In contrast, node t is allowed to use its knowledge of all messages $(a_{t',\mathcal{T}}^q)_{t' \in \mathcal{T}}$ during its decoding. Notice that this set of strategies in particular includes XOR-based inter-message coding as proposed in [1], [12] where each XOR is independently multicast to a set of receivers that can remove the non-desired parts in the XOR.

Theorem 2 (Converse under Non-Cooperative Transmissions). *Under combinatorial file assignments and non-cooperative transmissions of sub-IVAs, the achievable NDT-computation tradeoff is bounded as:*

$$\Delta(r) \geq \frac{1}{K} \left(1 - \frac{r}{K}\right) \cdot \frac{(K-2)r + K - 1}{r(K-1)}. \quad (13)$$

Proof: The converse proof for the r -fold cooperative channel with message-splitting approach is given in Section V. The tradeoff can be obtain by applying Lemma 1. ■

Then, we obtain the following corollary.

Corollary 1. *For $K \geq 5$ and $r = \lfloor (K-1)/2 \rfloor$ the bound in (12) is strictly smaller than the bound in (13).*

Proof: For K even and $r = \lfloor (K-1)/2 \rfloor = (K-2)/2$, the expression in (13) evaluates to:

$$\frac{1}{K} \left(1 - \frac{r}{K}\right) \cdot \frac{(K-2)^2 + 2(K-1)}{(K-2)(K-1)}. \quad (14)$$

which is strictly larger than the achievable bound in (12) for all even $K \geq 6$. Similarly, for K odd and $r = \lfloor (K-1)/2 \rfloor = (K-1)/2$, the bound in (13) simplifies to

$$\frac{1}{K(K-1)} \left(1 - \frac{r}{K}\right). \quad (15)$$

The obtained bound is again larger than the achievable bound in (12) for all $K \geq 3$. ■

IV. THE NEW IA & ZF SCHEME

A. A scheme for $r = \lfloor (K-1)/2 \rfloor$

Each message is cooperatively transmitted by a size- r set of nodes \mathcal{T} so that it is received at a given node $k \in [K] \setminus \mathcal{T}$ while zero-forced at a group \mathcal{S} of $r-1$ nodes in $[K] \setminus \{\mathcal{T} \cup \{k\}\}$. There is a set $\mathcal{L} \subset [K] \setminus \{\mathcal{S} \cup \mathcal{T} \cup \{k\}\}$ of $K - (r+r-1+1) = K - 2r$ nodes where the signal is experienced as interference. We choose to precode all messages that cause interference at a given set of Nodes \mathcal{L} by the same precoding matrix $\mathbf{U}_{\mathcal{L}}$, for $\mathcal{L} \in [K]$, with $|\mathcal{L}| = K - 2r$. We want to construct the precoding matrix $\mathbf{U}_{\mathcal{L}}$ so that all interferences at the same set of nodes \mathcal{L} align, thus leaving the remaining space for signaling dimensions. To summarize, if we use precoding matrix $\mathbf{U}_{\mathcal{L}}$ for the transmission of a message from group \mathcal{T} to Node k , then we zero-force this signal at all nodes in $[K] \setminus \{\mathcal{T} \cup \mathcal{L} \cup \{k\}\}$. We ensure that this signal is aligned with all other interference signals precoded with $\mathbf{U}_{\mathcal{L}}$ at the nodes $\ell \in \mathcal{L}$, which are the only nodes where it causes interference.

Based on the above design principle, the key design choice in our scheme is the assignment of a precoding matrix to each message. Equivalently, for each message, we select the set of nodes \mathcal{L} at which this message causes interference, and assign the corresponding precoding matrix $\mathbf{U}_{\mathcal{L}}$. This assignment fully determines both the interference alignment structure and the zero-forcing constraints of the scheme.

For the case $K \geq 5$, the assignment of precoding matrices follows a cyclic construction, as summarized in Algorithm 1.

Notice that the size of \mathcal{L} is $K - 2r$, so we select exactly r different sets \mathcal{T} in column k when K is even. We select $2r - 1 = K - 2$ different sets \mathcal{T} when K is odd.

TABLE I
CHOICE OF PRECODING MATRICES IN OUR SCHEME FOR $K = 5$ AND $r = 2$. EACH SIGNAL THAT IS PRECODED WITH MATRIX \mathbf{U}_5 IS ZERO-FORCED AT THE UNIQUE NODE $i \in [K] \setminus (\mathcal{T} \cup \{k, 5\})$.

$\mathcal{T} \setminus k$	1	2	3	4	5
{1,2}	x	x	\mathbf{U}_5	\mathbf{U}_5	o
{1,3}	x	\mathbf{U}_5	x	\mathbf{U}_5	o
{1,4}	x	\mathbf{U}_5	\mathbf{U}_5	x	o
{1,5}	x	o	o	o	x
{2,3}	\mathbf{U}_5	x	x	\mathbf{U}_5	o
{2,4}	\mathbf{U}_5	x	\mathbf{U}_5	x	o
{2,5}	o	x	o	o	x
{3,4}	\mathbf{U}_5	\mathbf{U}_5	x	x	o
{3,5}	o	o	x	o	x
{4,5}	o	o	o	x	x

For example, as depicted in Table I, for $K = 5$, $r = 2$, $\mathcal{L} = \{5\}$ and $k = 2$, the following sets \mathcal{T} send a message to Node 2 using precoding matrix \mathbf{U}_5 :

$$\mathcal{T} \in \{\{1, t\} : t \in [5] \setminus \{1, 2, 5\}\} = \{\{1, 3\}, \{1, 4\}\}, \quad (16)$$

where here we associated the index 1 with $k-1$ and t runs over the remaining set $[5] \setminus \{1, 2, 5\}$. Since K is odd, an additional assignment is required according to Algorithm 1. Specifically, the precoding matrix \mathbf{U}_5 must also be assigned to the entry corresponding to $\mathcal{T} = \{1, 4\}$ and $k' = k + 1 = 3$, as the transmit set $\{1, 4\}$ does not include 3. For clarity, we focus on the assignment of \mathbf{U}_5 with $k = 2$. The analysis for other precoding matrices and columns follows analogously by symmetry.

Another example is given in [13] for $K = 6$, $r = 2$.

B. Analysis of SDoF

If all channel coefficients after zero-forcing that multiply the same precoding matrix are algebraically independent, an IA coding scheme can be constructed by following the framework in [7] and applying the precoding matrix assignment method in Algorithm 1. For $K = 5$, the desired algebraic independence is proved analytically in the arXiv version [13]. For larger values of K however, the analysis becomes considerably more complicated and we verified algebraic independence directly by computing the Jacobian matrix using MATLAB². This verification has been carried out for all $K \in \{6, \dots, 15\}$. This range is chosen due to the computational complexity of the symbolic Jacobian computation.

To evaluate the SDoF achieved by our scheme, we notice that the DoF per node is given by the ratio between the dimension of the desired signal subspace and the dimension of the desired signal subspace and interference subspace.

²The code is at: <https://github.com/yue-bi/wdc-jacobian-verification>.

Algorithm 1 Assignment of precoding matrices for $K \geq 5$

Throughout Algorithm 1, all indices are interpreted according to the natural cyclic order of the set $[K] \setminus \mathcal{L}$. In particular, the notions of successor (e.g., $k + 1$) and consecutive indices are defined according to this cyclic order, rather than the natural integer ordering.

Require: Number of nodes K , cooperation parameter r , index set \mathcal{L}

- 1: **for all** $k \in [K] \setminus \mathcal{L}$ **do**
 - 2: Fix $r - 1$ consecutive nodes $\{k - r + 1, \dots, k - 1\} \subseteq [K] \setminus \mathcal{L}$.
 - 3: **for all** $t \in [K] \setminus (\{k - r + 1, \dots, k\} \cup \mathcal{L})$ **do**
 - 4: Form the set $\mathcal{T} \leftarrow \{k - r + 1, \dots, k - 1, t\}$
 - 5: Assign $\mathbf{U}_{\mathcal{L}}$ to row \mathcal{T} and column k
 - 6: **if** K is odd **and** $k + 1 \notin \mathcal{T}$ **then**
 - 7: Assign $\mathbf{U}_{\mathcal{L}}$ to row \mathcal{T} and column $k + 1$
 - 8: **end if**
 - 9: **end for**
 - 10: **end for**
-

For sufficiently large message lengths, the ratio between the dimensions of the desired signal subspace and the interference subspace asymptotically equals the ratio between the number of codewords intended for a given Node and the number of distinct precoding matrices in the interference space.

As mentioned earlier, each column is filled with r matrices when K is even, and with $2r - 1 = K - 2$ matrices when K is odd. This implies that for a node $k \notin \mathcal{L}$, the precoding matrix $\mathbf{U}_{\mathcal{L}}$ encodes r messages intended to Node k when K is even, and it encodes $K - 2$ messages intended to Node k when K is odd. Since $|\mathcal{L}| = K - 2r$, the number of precoding matrices that encode a message intended for Node k equals $\binom{K-1}{K-2r}$. Meanwhile, Node k is only interfered by codewords that are premultiplied by $\mathbf{U}_{\mathcal{L}}$ with $k \in \mathcal{L}$. The number of interfering precoding matrices is $\binom{K-1}{K-2r-1}$. Therefore, for $r = \lfloor (K - 1)/2 \rfloor$, the following DoF is achievable to each Node k when we let the length of message tends to infinity:

$$\begin{cases} \frac{\binom{K-1}{K-2r} r}{\binom{K-1}{K-2r} r + \binom{K-1}{K-2r-1}} = \frac{(K-2)^2}{(K-2)^2 + 4}, & K \text{ even,} \\ \frac{\binom{K-1}{K-2r} (K-2)}{\binom{K-1}{K-2r} (K-2) + \binom{K-1}{K-2r-1}} = \frac{(K-1)(K-2)}{(K-1)(K-2) + 1}, & K \text{ odd.} \end{cases} \quad (17)$$

The achievable SDoF can be obtained by multiplying the above results by K , as the scheme is symmetric.

V. PROOF OF THE CONVERSE RESULT (THEOREM 2)

We recall that the converse proof is derived under the assumption that each IVA $a_{q,p}$ is split into r equal-length sub-messages $(a_{t,\mathcal{T}}^q)_{t \in \mathcal{T}}$ and $\mathcal{T} = \{k : p \in \mathcal{M}_k\}$. Denote the set of all sub-IVAs by

$$\mathcal{V}_{\text{total}} := \{a_{t,\mathcal{T}}^q : \mathcal{T} \subset [[K]]^r, t \in \mathcal{T}, q \in [K] \setminus \mathcal{T}\}. \quad (18)$$

We have the following lemma.

Lemma 2. Fix $P > 0$ and a rate tuple in $\mathcal{C}(P)$. For any indices $j, t \in [K]$ with $j \neq t$, the following high-SNR inequality holds for R_a the rate of message a :

$$\overline{\lim}_{P \rightarrow \infty} \frac{1}{\log P} \sum_{a \in \mathcal{V}} R_a \leq 1, \quad (19)$$

where \mathcal{V} is defined as

$$\mathcal{V} \triangleq \mathcal{V}^r \cup \mathcal{V}^t, \quad (20)$$

for

$$\mathcal{V}^{\text{Rx}} \triangleq \{a_{u,\mathcal{T}}^j : \mathcal{T} \in [[K] \setminus \{j\}]^r \text{ and } u \in \mathcal{T}\}, \quad (21)$$

and

$$\mathcal{V}^{\text{Tx}} \triangleq \bigcup_{i \in [K] \setminus \{j,t\}} \mathcal{V}^{\text{Tx},i}, \quad (22)$$

$$\mathcal{V}^{\text{Tx},i} \triangleq \{a_{t,\mathcal{T}}^i : \mathcal{T} \in [[K] \setminus \{j, \dots, i\}]^r \text{ and } t \in \mathcal{T}\}. \quad (23)$$

In all the above expressions, the index set $\{j, \dots, i\}$ needs to be interpreted according to the natural cyclic order of the set $[K] \setminus \{t\}$.

Proof. Without loss of generality, we choose $j = 2$ and $t = 1$ throughout this proof.

Fix $P > 0$ and any rate tuple in $\mathcal{C}(P)$. Consider encoding and decoding functions $\{f_q^{(n)}\}$ and $\{g^{(n)}\}$ such that $p^{(n)}(\text{error}) \rightarrow 0$ as $n \rightarrow \infty$. Fix a blocklength n and set $\mathcal{F} \triangleq \{3, \dots, K\}$. Let \mathcal{H} denote the collection of all channel coefficients. For any $\mathcal{S} \subseteq [K]$, define $\mathbf{Y}_{\mathcal{S}} \triangleq (\mathbf{Y}_p : p \in \mathcal{S})$ and $\mathbf{Z}_{\mathcal{S}} \triangleq (\mathbf{Z}_p : p \in \mathcal{S})$. All other messages are collectively denoted by

$$\bar{\mathcal{V}} \triangleq \mathcal{V}_{\text{total}} \setminus \mathcal{V}. \quad (24)$$

By independence of all messages, channel coefficients, and noise sequences,

$$\begin{aligned} H(\mathcal{V}) &= H(\mathcal{V} | \bar{\mathcal{V}}, \mathcal{H}) \\ &= I(\mathcal{V}; \mathbf{Y}_2 | \bar{\mathcal{V}}, \mathcal{H}) + H(\mathcal{V} | \mathbf{Y}_2, \bar{\mathcal{V}}, \mathcal{H}) \\ &= h(\mathbf{Y}_2 | \bar{\mathcal{V}}, \mathcal{H}) - h(\mathbf{Z}_2) + H(\mathcal{V} | \mathbf{Y}_2, \bar{\mathcal{V}}, \mathcal{H}) \\ &\leq n \cdot \log(1 + (K-1)PH_{\text{max}}^2) + H(\mathcal{V} | \mathbf{Y}_2, \bar{\mathcal{V}}, \mathcal{H}). \end{aligned} \quad (25)$$

In the following we show that the last entropy term grows slower than $n \log P$, which will conclude the proof of the lemma. To this end, start by decomposing the entropy term as

$$\begin{aligned} H(\mathcal{V} | \mathbf{Y}_2, \bar{\mathcal{V}}, \mathcal{H}) &= H(\mathcal{V}^{\text{Rx}} | \mathbf{Y}_2, \bar{\mathcal{V}}, \mathcal{H}) + H(\mathcal{V}^{\text{Tx}} | \mathbf{Y}_2, \bar{\mathcal{V}}, \mathcal{V}^{\text{Rx}}, \mathcal{H}). \end{aligned} \quad (26)$$

The first term grows sublinearly in n by Fano's inequality because the messages in \mathcal{V}^{Rx} need to be decoded from \mathbf{Y}_2 :

$$H(\mathcal{V}^{\text{Rx}} | \mathbf{Y}_2, \bar{\mathcal{V}}, \mathcal{H}) \leq n \epsilon_n, \quad (27)$$

for some $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$.

For the second term, we write

$$\begin{aligned} H(\mathcal{V}^{\text{Tx}} | \mathbf{Y}_2, \bar{\mathcal{V}}, \mathcal{V}^{\text{Rx}}, \mathcal{H}) &= H(\mathcal{V}^{\text{Tx}} | \mathbf{Y}_2, \mathbf{Y}_{\mathcal{F}}, \bar{\mathcal{V}}, \mathcal{V}^{\text{Rx}}, \mathcal{H}) \\ &\quad + I(\mathcal{V}^{\text{Tx}}; \mathbf{Y}_{\mathcal{F}} | \mathcal{V}^{\text{Rx}}, \mathbf{Y}_2, \bar{\mathcal{V}}, \mathcal{H}), \end{aligned} \quad (28)$$

and observe that

$$\begin{aligned} I(\mathcal{V}^{\text{Tx}}; \mathbf{Y}_{\mathcal{F}} | \mathcal{V}^{\text{Rx}}, \mathbf{Y}_2, \bar{\mathcal{V}}, \mathcal{H}) &\stackrel{(a)}{=} h(\tilde{\mathbf{Y}}_{\mathcal{F}} | \mathcal{V}^{\text{Rx}}, \tilde{\mathbf{Y}}_2, \bar{\mathcal{V}}, \mathcal{H}) - h(\mathbf{Z}_{\mathcal{F}}) \end{aligned} \quad (29)$$

$$\stackrel{(b)}{=} h(\tilde{\mathbf{Y}}_{\mathcal{F}} | \mathcal{V}^{\text{Rx}}, \tilde{\mathbf{Y}}_2, \bar{\mathcal{V}}, \mathcal{H}, E=1) - h(\mathbf{Z}_{\mathcal{F}}) \quad (30)$$

$$\leq h(\mathbf{Z}'_{\mathcal{F}} | \mathcal{V}^{\text{Rx}}, \tilde{\mathbf{Y}}_2, \bar{\mathcal{V}}, \mathcal{H}, E=1) - h(\mathbf{Z}_{\mathcal{F}}) \quad (31)$$

$$\leq h(\mathbf{Z}'_{\mathcal{F}}) - h(\mathbf{Z}_{\mathcal{F}}) + n \bar{\epsilon}_n, \quad (32)$$

where $\bar{\epsilon}_n$ is a sequence that vanishes in n and the (in)equalities are justified as follows:

- In (a), we define $\tilde{\mathbf{Y}}_p \triangleq \mathbf{H}_{p,t} \mathbf{X}_t + \mathbf{Z}_p$, $p \in [K]$. and $\tilde{\mathbf{Y}}_{\mathcal{S}} \triangleq (\tilde{\mathbf{Y}}_p : p \in \mathcal{S})$. Since \mathcal{V}^{Tx} is carried only by the input of Node t , conditioning on $(\mathcal{V}^{\text{Rx}}, \bar{\mathcal{V}}, \mathcal{H})$ removes all other signal components, yielding the equality in (a).
- In (b), we introduce the binary random variable E , which equals 1 if for all $t' \in [n]$, the channel coefficient $H_{j,t}(t') \neq 0$. Notice that $\Pr[E=1] = 1$ since the channel coefficients are drawn independently from continuous distributions.
- To see (c), define for each $p \in \mathcal{F}$ the noise-modified variables $\mathbf{Z}'_p \triangleq \mathbf{Z}_p - \mathbf{H}_{p,t} \mathbf{H}_{j,t}^{-1} \mathbf{Z}_j$, and $\mathbf{Z}'_{\mathcal{F}} \triangleq (\mathbf{Z}'_p : p \in \mathcal{F})$. Notice that in the event $E=1$, the signal component of

$\tilde{\mathbf{Y}}_{\mathcal{F}}$ can be perfectly canceled, leaving only the effective noise $\mathbf{Z}'_{\mathcal{F}}$.

We finally show

$$\begin{aligned} H(\mathcal{V}^{\text{Tx}} | \mathbf{Y}_2, \mathbf{Y}_{\mathcal{F}}, \mathcal{V}^{\text{Rx}}, \bar{\mathcal{V}}, \mathcal{H}) &= \sum_{i=3}^K H(\mathcal{V}^{\text{Tx},i} | \mathbf{Y}_2, \mathbf{Y}_{\mathcal{F}}, \mathcal{V}^{\text{Rx}}, \bar{\mathcal{V}}, \mathcal{V}^{\text{Tx},3}, \dots, \mathcal{V}^{\text{Tx},i-1}, \mathcal{H}) \end{aligned} \quad (33)$$

$$\leq \sum_{i=3}^K \epsilon_{i,n} n, \quad (34)$$

where $\epsilon_{i,n}$ denotes a vanishing sequence for each i and the last inequality holds by the following considerations.

For example, notice that the sub-IVAs known to node 3, i.e., $\{a_{p,\mathcal{T}}^q : 3 \in \mathcal{T}\}$ are all contained in the set $(\mathcal{V}^{\text{Rx}} \cup \bar{\mathcal{V}})$. Moreover, $\mathcal{V}^{\text{Tx},3}$ only contains sub-IVAs intended for node 3. Since $\mathbf{Y}_{\mathcal{F}}$ contains \mathbf{Y}_3 , we thus have by Fano's Inequality equality:

$$H(\mathcal{V}^{\text{Tx},3} | \mathbf{Y}_2, \mathbf{Y}_{\mathcal{F}}, \mathcal{V}^{\text{Rx}}, \bar{\mathcal{V}}) \leq \epsilon_{3,n} n, \quad (35)$$

for a sequence $\epsilon_{3,n} \rightarrow 0$ as $n \rightarrow \infty$.

For each $i = 4, \dots, K$, the sub-IVAs known to Node i , i.e., $\{a_{p,\mathcal{T}}^q : i \in \mathcal{T}\}$, are contained in the set $(\mathcal{V}^{\text{Rx}} \cup \bar{\mathcal{V}} \cup \mathcal{V}^{\text{Tx},3} \cup \dots \cup \mathcal{V}^{\text{Tx},i-1})$. Since $\mathcal{V}^{\text{Tx},i}$ only contains sub-IVAs intended for node i and \mathbf{Y}_i is contained in $\mathbf{Y}_{\mathcal{F}}$, we have by Fano's inequality:

$$H(\mathcal{V}^{\text{Tx},i} | \mathbf{Y}_2, \mathbf{Y}_{\mathcal{F}}, \mathcal{V}^{\text{Rx}}, \bar{\mathcal{V}}, \mathcal{V}^{\text{Tx},3}, \dots, \mathcal{V}^{\text{Tx},i-1}) \leq \epsilon_{i,n} n, \quad (36)$$

for a sequence $\epsilon_{i,n} \rightarrow 0$ as $n \rightarrow \infty$. This proves (34). Combining (25)–(34), we obtain

$$\begin{aligned} H(\mathcal{V}) &\leq n \cdot \log(1 + (K-1)PH_{\text{max}}^2) + n \epsilon_n \\ &\quad + h(\mathbf{Z}'_{\mathcal{F}}) - h(\mathbf{Z}_{\mathcal{F}}) + n \bar{\epsilon}_n. \end{aligned} \quad (37)$$

Dividing by $n \log P$ and letting $n \rightarrow \infty$ and $P \rightarrow \infty$ establishes (19). \square

The proof is obtained by summing the upper bound in Lemma 2 over all index pairs $(t, j) \in [K]$ with $t \neq j$, and then normalizing. We obtain $\text{SDoF} \leq |\mathcal{V}_{\text{total}}|/|\mathcal{V}|$, thus

$$\text{SDoF} \leq \frac{\binom{K}{r} (K-r)}{\binom{K-2}{r} + \binom{K-1}{r} \cdot r} = \frac{K(K-1)r}{(K-2)r + K-1}. \quad (38)$$

VI. CONCLUSION

We propose a new wireless MapReduce coding scheme based on IA and ZF for computation load $r = \lfloor (K-1)/2 \rfloor$, which strictly improves upon all previous achievable NDTs. We also establish a converse for non-cooperative schemes, where IVAs are split into submessages and transmitted by single nodes. The results show that such schemes are strictly suboptimal compared to our IA–ZF approach.

ACKNOWLEDGEMENTS

The work of the authors has been supported by the ERC under grant agreement 101125691.

REFERENCES

- [1] S. Li, M. A. Maddah-Ali, Q. Yu, and A. S. Avestimehr, "A fundamental tradeoff between computation and communication in distributed computing," *IEEE Transactions on Information Theory*, vol. 64, no. 1, pp. 109–128, Jan. 2018.
- [2] Q. Yan, S. Yang, and M. Wigger, "Storage-computation-communication tradeoff in distributed computing: Fundamental limits and complexity," *IEEE Transactions on Information Theory*, vol. 68, no. 8, pp. 5496–5512, 2022.
- [3] F. Xu, S. Shao, and M. Tao, "New results on the computation-communication tradeoff for heterogeneous coded distributed computing," *IEEE Transactions on Communications*, vol. 69, no. 4, pp. 2254–2270, Apr. 2021.
- [4] K. Wan, H. Sun, M. Ji, and G. Caire, "Distributed linearly separable computation," *IEEE Transactions on Information Theory*, vol. 68, no. 2, pp. 1259–1278, 2022.
- [5] F. Li, J. Chen, and Z. Wang, "Wireless mapreduce distributed computing," *IEEE Transactions on Information Theory*, vol. 65, no. 10, pp. 6101–6114, Oct. 2019.
- [6] Y. Bi, P. Ciblat, M. Wigger, and Y. Wu, "DoF of a cooperative X-channel with an application to distributed computing," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Espoo, Finland, Jun. 2022, pp. 566–571.
- [7] Y. Bi, M. Wigger, and Y. Wu, "Normalized delivery time of wireless mapreduce," *IEEE Transactions on Information Theory*, vol. 70, no. 10, pp. 7005–7022, 2024.
- [8] K. Yuan and Y. Wu, "Coded wireless distributed computing via interference alignment," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Espoo, Finland, 2022.
- [9] V. R. Cadambe and S. A. Jafar, "Interference alignment and degrees of freedom of the K-user interference channel," *IEEE Transactions on Information Theory*, vol. 54, no. 8, pp. 3425–3441, Aug. 2008.
- [10] S. Jafar and S. Shamai, "Degrees of freedom region of the MIMO X-channel," *IEEE Transactions on Information Theory*, vol. 54, no. 1, pp. 151–170, Jan. 2008.
- [11] V. S. Annapureddy, A. El Gamal, and V. V. Veeravalli, "Degrees of freedom of interference channels with CoMP transmission and reception," *IEEE Transactions on Information Theory*, vol. 58, no. 9, pp. 5740–5760, Sep. 2012.
- [12] J. Hachem, U. Niesen, and S. N. Diggavi, "Degrees of freedom of cache-aided wireless interference networks," *IEEE Transactions on Information Theory*, vol. 64, no. 7, pp. 5359–5380, Jul. 2018.
- [13] Y. Bi and M. Wigger, "Necessity of cooperative transmissions for wireless mapreduce," 2026. [Online]. Available: <https://arxiv.org/abs/2601.11844>