

Degrees of Freedom of Cache-Aided Wireless Interference Networks

Jad Hachem, Urs Niesen, and Suhas Diggavi

Abstract

We study the role of caches in wireless interference networks. We focus on content caching and delivery across a Gaussian interference network, where both transmitters and receivers are equipped with caches. We provide a constant-factor approximation of the system's degrees of freedom (DoF), for arbitrary number of transmitters, number of receivers, content library size, receiver cache size, and transmitter cache size (as long as the transmitters combined can store the entire content library among them). We demonstrate approximate optimality with respect to information-theoretic bounds that do not impose any restrictions on the caching and delivery strategies. Our characterization reveals three key insights. First, the approximate DoF is achieved using a strategy that separates the physical and network layers. This separation architecture is thus approximately optimal. Second, we show that increasing transmitter cache memory beyond what is needed to exactly store the entire library between all transmitters does not provide more than a constant-factor benefit to the DoF. A consequence is that transmit zero-forcing is not needed for approximate optimality. Third, we derive an interesting trade-off between the receiver memory and the number of transmitters needed for approximately maximal performance. In particular, if each receiver can store a constant fraction of the content library, then only a constant number of transmitters are needed. Solving the caching problem required formulating and solving a new communication problem, the symmetric multiple multicast X-channel, for which we provide an exact DoF characterization.

I. INTRODUCTION

Traditional communication networks focus on establishing a reliable connection between two fixed network nodes and are therefore connection centric. With the recent explosion in multimedia content, network usage has undergone a significant shift: users now want access to some specific content, regardless of its location in the network. Consequently, network architectures are shifting towards being content centric. These content-centric architectures make heavy use of in-network caching and, in doing so, redesign the protocol stack from the network layer upwards [1].

A natural question to ask is how the availability of in-network caches can be combined with the wireless physical layer and specifically with two fundamental properties of wireless communication: the broadcast and the superposition of transmitted signals. Recent work in the information theory literature has demonstrated that this combination can yield significant benefits. This information-theoretic approach to caching was introduced in the context of the noiseless broadcast channel in [2], where it was shown that significant performance gains can be obtained using cache memories at the *receivers*. In [3], the noiseless broadcast setting was extended to the interference channel, which is the simplest multiple-unicast wireless topology capturing both broadcast and superposition. The authors presented an achievable scheme showing performance gains using cache memories at the *transmitters*.

In this paper, we continue the study of the cache-aided wireless interference network, but we allow for caches at both the *transmitters and receivers* as shown in Fig. 1. Our main result (Theorem 1, Section III) is a complete constant-factor approximation of the degrees of freedom (DoF) of this network. The result is general, in that it holds for any number of transmitters and receivers, size of content library, transmitter cache size (large enough to collectively hold the entire content library), and receiver cache size. Moreover,

J. Hachem and S. Diggavi are with the Department of Electrical Engineering, University of California, Los Angeles. Emails: jadhachem@ucla.edu, suhas@ee.ucla.edu.

U. Niesen is with the Qualcomm NJ Research Center, Bridgewater, NJ. Email: urs.niesen@ieee.org.

This work was supported in part by NSF grant #1423271.

This paper was presented in part at the 2016 IEEE International Symposium on Information Theory in Barcelona, Spain.

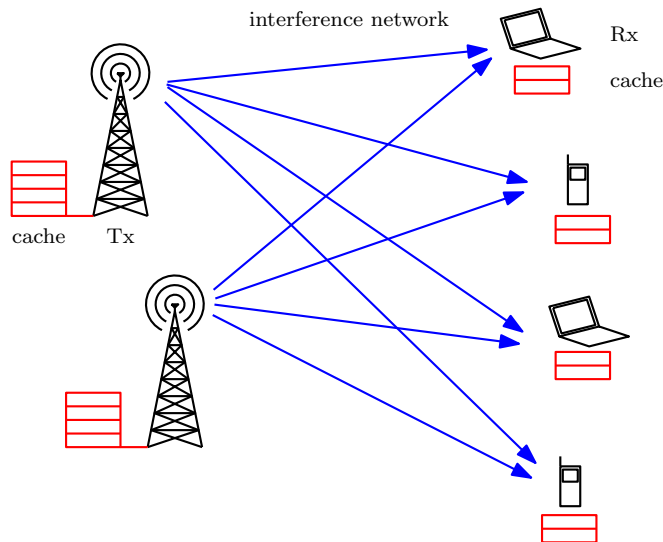


Fig. 1. Caching in a wireless interference network. Caches (in red) are placed at all network nodes.

our converse holds for arbitrary caching and transmission functions, and imposes no restrictions as done in prior work.

Several architectural and design insights emerge from this degrees-of-freedom approximation.

- 1) Our achievable scheme introduces a novel separation of the physical and network layers, thus redesigning the protocol stack from the network layer downwards. From the order-wise matching converse, we hence see that this separation is approximately optimal.
- 2) Once the transmitter caches are large enough to collectively hold the entire content library, increasing the transmitter memory further can lead to at most a constant-factor improvement in the system's degrees of freedom. In particular, and perhaps surprisingly, this implies that transmit zero-forcing is not needed for approximately optimal performance.
- 3) There is a trade-off between the number of transmitters needed for (approximately) maximal system performance and the amount of receiver cache memory. As the receiver memory increases, the required number of transmitters decreases, down to a constant when the memory is a constant fraction of the entire content library.

There are three seemingly natural network-layer abstractions for this problem. The first network-layer abstraction treats the physical layer as a standard interference channel and transforms it into non-interacting bit pipes between disjoint transmitter-receiver pairs. This approach is inefficient. The second network-layer abstraction treats the physical layer as an X-channel and transforms it into non-interacting bit pipes between each transmitter and each receiver. The third network-layer abstraction treats the physical layer as multiple broadcast channels: it creates a broadcast link from each transmitter to all receivers. The last two approaches turn out to be approximately optimal in special circumstances: the second when the receivers have no memory, and the third when they have enough memory to each store almost all the content library. In this paper, we propose a network-layer abstraction that creates X-channel *multicast* bit pipes, each sent by a transmitter and intended for a subset of receivers whose size depends on the receiver memory. This abstraction generalizes the above two approaches, and we show that it is in fact order-optimal for all values of receiver memory.

Our solution to this problem requires solving a new communication problem at the physical layer that arises from the proposed separation architecture. This problem generalizes the X-channel setting studied in [4] by considering multiple *multicast* messages instead of just unicast. We focus on the symmetric case and provide a complete and exact DoF characterization of this symmetric multiple multicast X-channel problem, by proposing a strategy based on interference alignment and proving its optimality (see Theorem 2, Section IV).

Related Work

Content caching has a rich history and has been studied extensively, see for example [5] and references therein. Recent interest in content caching is motivated by Video-on-Demand systems for which efficient content placement and delivery schemes have been proposed in [6]–[9]. The impact of content popularity distributions on caching schemes has also been widely investigated, see for example [10]–[12]. Most of the literature has focused on wired networks, and the solutions there do not carry directly to wireless networks.

The information-theoretic framework for coded caching was introduced in [2] in the context of the deterministic broadcast channel. This has been extended to online caching systems [13], systems with delay-sensitive content [14], heterogeneous cache sizes [15], unequal file sizes [16], and improved converse arguments [17], [18]. Content caching and delivery in device-to-device networks, multi-server topologies, and heterogeneous wireless networks have been studied in [19]–[23]. This framework was also applied to hierarchical (tree) topologies in [24], and to non-uniform content popularities in [22], [23], [25]–[28]. Other related work includes [29], which derives scaling laws for content replication in multihop wireless networks, and [30], which explores distributed caching in mobile networks using device-to-device communications. Wireless caching networks with rate-limited fronthaul were explored in [31]. The benefit of coded caching when the caches are randomly distributed was studied in [32], and the benefits of adaptive content placement using knowledge of user requests were explored in [33].

More recently, this information-theoretic framework for coded caching has been extended in [3] to interference channels with caches at only the transmitters. The analysis in [3] focuses on three transmitters and three receivers; extensions to arbitrary numbers of transmitters and receivers were considered in [31]. Interference channels with caches *both* at transmitters and at receivers were considered in [34]–[36], all of which have a setup similar to the one in this paper. However, each of these three works has some restrictions on the setup. The authors in [34] focus on one-shot linear schemes, while [35] prohibits inter-file coding during placement and limits the number of receivers to three. Our prior work [36] studies the same setup but with only two transmitters and two receivers. The work in this paper differs from those above in that it considers an arbitrary number of transmitters and receivers and proves order-optimality using outer bounds that assume *no restrictions* on the scheme. We note that this work was developed independently from [35].

Organization

The remainder of this paper is organized as follows. Section II introduces the problem setting and establishes notation. Section III states the paper’s main results. Section IV presents the separation architecture in detail; Section V gives the interference alignment strategy used at the physical layer. Section VI proves the order-optimality of our strategy. Section VII discusses possible extensions to the separation architecture. We defer additional proofs to the appendices.

II. PROBLEM SETTING

We consider the setup shown in Fig. 2. A content library contains N files W_1, \dots, W_N of size F bits each. A total of K_r users will each request one of these files, which must be transmitted across a $K_t \times K_r$ time-varying Gaussian interference channel whose receivers are the system’s users. We will hence use the terms “receiver” and “user” interchangeably. Our goal is to reliably transmit these files to the users with the help of caches at both the transmitters and the receivers.

The system operates in two phases, a *placement phase* and a *delivery phase*. In the placement phase, the transmitter and receiver caches are filled as an arbitrary function of the content library. The transmitter caches are able to store $M_t F$ bits; the receiver caches are able to store $M_r F$ bits. We refer to M_t and M_r as the transmitter and receiver cache sizes, respectively. Other than the memory constraints, we impose no restrictions on the caching functions (in particular, we allow the caches to arbitrarily code across files). In

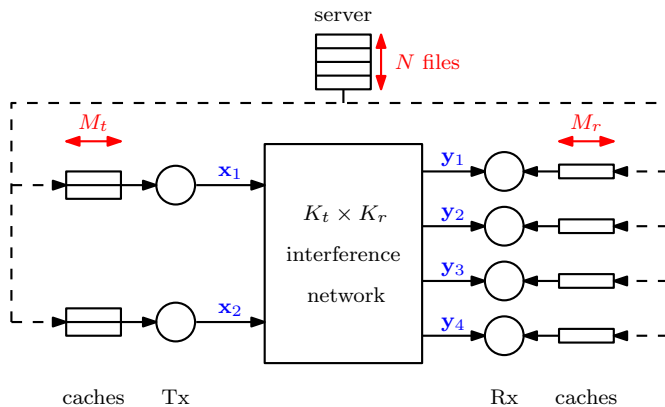


Fig. 2. The caching problem, with $K_t = 2$ transmitters and $K_r = 4$ receivers. The server holds a content library of N files. Information about these files is placed in the transmitter caches of size M_t and in the receiver caches of size M_r during a placement phase (indicated by dashed lines). During the subsequent delivery phase (indicated by solid lines), each user requests one file, and all the requested files have to be delivered over the interference network.

this paper, we consider all values of $M_r \geq 0$, but we restrict ourselves to the case where the transmitter caches can collectively store the entire content library,¹ i.e.,

$$M_t \geq N/K_t. \quad (1)$$

The delivery phase takes place after the placement phase is completed. In the beginning of the delivery phase, each user requests one of the N files. We denote by $\mathbf{u} = (u_1, \dots, u_{K_r})$ the vector of user demands, such that user i requests file W_{u_i} . These requests are communicated to the transmitters, and each transmitter j responds by sending a codeword $\mathbf{x}_j = (x_j(1), \dots, x_j(T))$ of block length T into the interference channel. We impose a power constraint over every channel input \mathbf{x}_j ,

$$\frac{1}{T} \|\mathbf{x}_j\|^2 \leq \text{SNR}, \quad \forall j = 1, \dots, K_t.$$

Note that each transmitter only has access to its own cache, so that \mathbf{x}_j only depends on the contents of transmitter j 's cache and the user requests \mathbf{u} . We impose no other constraint on the channel coding function (in particular, we explicitly allow for coding across time using potentially nonlinear schemes).

Receiver i observes a noisy linear combination of all the transmitted codewords,

$$y_i(\tau) \triangleq \sum_{j=1}^{K_t} h_{ij}(\tau) x_j(\tau) + z_i(\tau),$$

for all time instants $\tau = 1, \dots, T$, where $z_i(\tau)$ is iid unit-variance additive Gaussian noise, and $h_{ij}(\tau)$ are independent time-varying random channel coefficients obeying some continuous probability distribution. We can rewrite the channel outputs in vector form as

$$\mathbf{y}_i = \sum_{j=1}^{K_t} \mathbf{H}_{ij} \mathbf{x}_j + \mathbf{z}_i, \quad (2)$$

where \mathbf{H}_{ij} is a diagonal matrix representing the channel coefficients over the block length T .

For fixed values of M_t , M_r , and SNR, we say that a transmission rate $R = F/T$ is achievable if there exists a coding scheme such that all the users can decode their requested files with vanishing error probability. More formally, R is *achievable for demand vector* \mathbf{u} if

$$\max_{i \in \{1, \dots, K_r\}} \mathbb{P}(\hat{W}_i \neq W_{u_i}) \rightarrow 0 \text{ as } T \rightarrow \infty,$$

¹To achieve any positive DoF, the minimum requirement is that $K_t M_t + M_r \geq N$, i.e., that all the transmitter caches and any single receiver cache can collectively store the entire content library. We impose the slightly stronger requirement $K_t M_t \geq N$ since it is the regime of most practical interest, and since it simplifies the analysis.




Fig. 3. Approximate reciprocal DoF of the 2×4 cache-aided interference network with 4 files as a function of receiver cache size M_r .

where \hat{W}_i indicates the reconstruction of file W_{u_i} by user i . Note that R is fixed as T , and hence F , go to infinity. We say R is *achievable* if it is achievable for all demand vectors \mathbf{u} .

We define the optimal transmission rate $R^*(\text{SNR})$ as the supremum of all achievable rates for a given SNR (and number of files, cache sizes, and number of transmitters/receivers). In the remainder of this paper we will focus on the degrees of freedom (DoF) defined as

$$\text{DoF} \triangleq \lim_{\text{SNR} \rightarrow \infty} \frac{R^*(\text{SNR})}{\frac{1}{2} \log \text{SNR}}. \quad (3)$$

We will often work with the reciprocal $1/\text{DoF}$ because it is a convex function of (M_t, M_r) .

III. MAIN RESULTS

The main result of this paper is a complete constant-factor approximation of the DoF for the cache-aided wireless interference network. In order to state the result, we define the function $d(N, K_t, K_r, M_t, M_r)$ through

$$\frac{1}{d(N, K_t, K_r, M_t, M_r)} \triangleq \frac{K_t - 1 + \min\{\frac{K_r}{\kappa+1}, N\}}{K_t} \cdot \left(1 - \frac{\kappa}{K_r}\right), \quad (4)$$

for any N, K_t, K_r, M_t , and $M_r = \kappa N / K_r$ with $\kappa \in \{0, 1, \dots, K_r\}$, and the lower convex envelope of these points for all other $M_r \in [0, N]$.

Theorem 1. *The degrees of freedom DoF of the $K_t \times K_r$ cache-aided interference network with N files, transmitter cache size $M_t \in [N/K_t, N]$, and receiver cache size $M_r \in [0, N]$ satisfies*

$$d(N, K_t, K_r, M_t, M_r) \leq \text{DoF} \leq 13.5 \cdot d(N, K_t, K_r, M_t, M_r).$$

In terms of the rate $R^*(\text{SNR})$ of the system, Theorem 1 can be interpreted using (3) as

$$d(\cdot) \cdot \frac{1}{2} \log \text{SNR} - o(\log \text{SNR}) \leq R^*(\text{SNR}) \leq 13.5 \cdot d(\cdot) \cdot \frac{1}{2} \log \text{SNR} + o(\log \text{SNR}),$$

when SNR grows, where we have used $d(\cdot)$ instead of $d(N, K_t, K_r, M_t, M_r)$ for simplicity.

The constant 13.5 in Theorem 1 is the result of some loosening of inequalities in order to simplify the analysis. We numerically observe that the multiplicative gap does not exceed 4.16 for $N, K_t, K_r \leq 100$.

Example 1. Consider a cache-aided interference network with $K_t = 2$ transmitters and $K_r = 4$ receivers. The number of files is $N = 4$, and the transmitter cache size is $M_t = 2$ (sufficient to exactly hold the entire content library collectively at the transmitters). As a function of the receiver cache size $M_r \in [0, 4]$, the reciprocal of the approximate degrees of freedom $d(\cdot) = d(N, K_t, K_r, M_t, \cdot)$ is shown in Fig. 3. \diamond

The coding scheme achieving the lower bound on DoF in Theorem 1 uses separate network and physical layers. The two layers interface using a set of multicast messages from each transmitter to many subsets of receivers. At the physical layer, an interference alignment scheme (generalizing the scheme from [4]) delivers these messages across the interference channel with vanishing error probability and at optimal degrees of freedom. At the network layer, a caching and delivery strategy generalizing the one in [2] is used to deliver the requested content to the users, utilizing the non-interacting error-free multicast bit pipes created by the physical layer. The matching upper bound in Theorem 1 shows that this separation approach is without loss of order optimality. This separation architecture is described in more detail in Section IV.

In order to better understand the behavior of the system, we decompose the approximation of the sum degrees of freedom $K_r \text{DoF}$ provided by Theorem 1 into three components, or gains. These are: an interference alignment (IA) gain g^{IA} , a local caching gain g^{LC} , and a global caching gain g^{GC} , forming

$$K_r \text{DoF} \approx K_r d(N, K_t, K_r, M_t, M_r) = \underbrace{\frac{K_t K_r}{K_t + K_r - 1}}_{g^{\text{IA}}} \cdot \underbrace{\frac{1}{1 - \frac{M_r}{N}}}_{g^{\text{LC}}} \cdot \underbrace{\frac{K_r M_r / N + 1}{\frac{M_r}{N} \left(\frac{1}{K_r} + \frac{1}{K_t - 1} \right)^{-1} + 1}}_{g^{\text{GC}}}. \quad (5)$$

We point out that, for ease of presentation, this decomposition is written for the case when the first term achieves the minimum in (4), i.e., $\frac{K_r}{\kappa+1} \leq N$. This includes the most relevant case when the content library N is larger than the number of receivers K_r . In fact, we focus on this case in most of the main body of the paper, particularly regarding the achievability and some of the intuition. A detailed discussion of the case $\frac{K_r}{\kappa+1} > N$, including a decomposition similar to (5), is given in Appendix A.

The term g^{IA} is the degrees of freedom achieved by communication using interference alignment and is the same as in the unicast X-channel problem [4]. It is the only gain present when the receiver cache size is zero. In other words, it is the baseline degrees of freedom without caching (see for example Fig. 3 when $M_r = 0$).

When the receiver cache size is non-zero, we get two improvements, in analogy to the two gains described in the broadcast caching setup in [2]. The local caching gain reflects that each user already has some information about the requested file locally in its cache. Hence, g^{LC} is a function of M_r/N , the fraction of each file stored in a single receiver cache. On the other hand, the global gain derives from the coding opportunities created by storing different content at different users, and from the multicast links created to serve coded information useful to many users at once. This gain depends on the total amount of receiver memory, as is reflected by the $K_r M_r / N$ term in the numerator of g^{GC} .

It is interesting to see how each of these gains scales with the various system parameters K_t , K_r , M_t/N , and M_r/N . In order to separate the different gains, we work with the logarithm

$$\log K_r \text{DoF} \approx \log g^{\text{IA}} + \log g^{\text{LC}} + \log g^{\text{GC}}$$

of the sum degrees of freedom. By varying the different parameters, we can plot how both the sum DoF and its individual components evolve.

1) *Scaling with transmitter memory M_t* : Notice in Theorem 1 that the DoF approximation does not involve the transmitter memory M_t . Thus, once $M_t = N/K_t$, just enough to store the entire content library between all transmitters, any increase in the transmit memory will only lead to at most a constant-factor improvement in the DoF.

The strategy used to achieve the lower bound in Theorem 1 (see Section IV for details) stores uncoded *nonoverlapping* file parts in each transmit cache. This is done regardless of the transmitter memory M_t and the receiver memory M_r . Since this is an order-optimal strategy, we conclude that the transmitters do not need to have any shared information. Consequently, and perhaps surprisingly, transmit zero-forcing is not needed for order-optimality and cannot provide more than a constant-factor DoF gain.

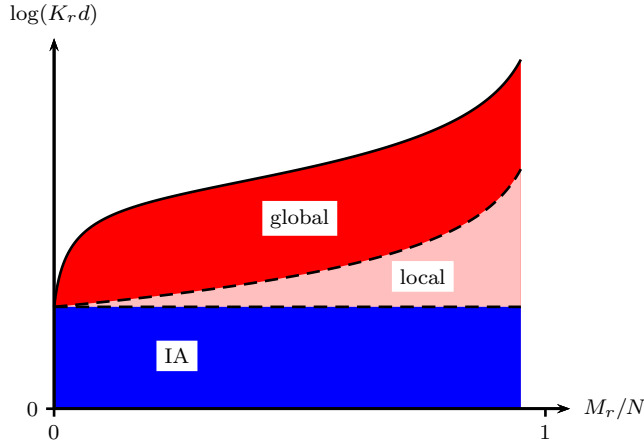


Fig. 4. DoF gains as a function of receiver cache size characterized by M_r/N .




Fig. 5. DoF gains as a function of number of receivers K_r .

2) *Scaling with receiver memory M_r* : Fig. 4 depicts the decomposition of the approximate sum degrees of freedom $K_r d \approx K_r \text{DoF}$ as a function of the receiver cache size M_r . As expected, the interference alignment gain g^{IA} does not depend on the receiver cache size and is hence constant. The local caching gain g^{LC} increases slowly with M_r and becomes relevant whenever each receiver can cache a significant fraction of the content library, say $M_r/N \geq 0.5$. The global caching gain g^{GC} increases much more quickly and is relevant whenever the cumulative receiver cache size is large, say $K_r M_r/N \geq 1$.

3) *Scaling with number of receivers K_r* : Fig. 5 depicts the decomposition of the approximate sum degrees of freedom $K_r d \approx K_r \text{DoF}$ as a function of the number of receivers K_r . The local caching gain g^{LC} is not a function of K_r and is hence constant as expected. In the limit as $K_r \rightarrow \infty$, the interference alignment gain g^{IA} converges to K_t . The global caching gain g^{GC} , on the other hand, behaves as

$$g^{\text{GC}} \approx \frac{K_r M_r/N + 1}{(K_t - 1)M_r/N + 1}$$

for large K_r . In particular, unlike the other two gains, the global gain does not converge to a limit and scales linearly with the number of receivers. Thus, for systems with larger number of receivers, the global caching gain becomes dominant.

4) *Scaling with number of transmitters K_t* : As the number of receivers K_r or the receive memory M_r increase, the sum DoF grows arbitrarily large. The same is not true as the number of transmitters K_t

increases. In fact, as $K_t \rightarrow \infty$, we find that $g^{\text{IA}} \rightarrow K_r$, $g^{\text{GC}} \rightarrow 1$, and the sum DoF converges to

$$\lim_{K_t \rightarrow \infty} K_r \text{DoF} \approx \frac{K_r}{1 - M_r/N}. \quad (6)$$

This is not surprising, since, with a large number of transmitters, interference alignment effectively creates K_r orthogonal links from each transmitter to the receivers, each of DoF approaching 1. With the absence of multicast due to these orthogonal links, the global caching gain vanishes and the only caching gain left is the local one.

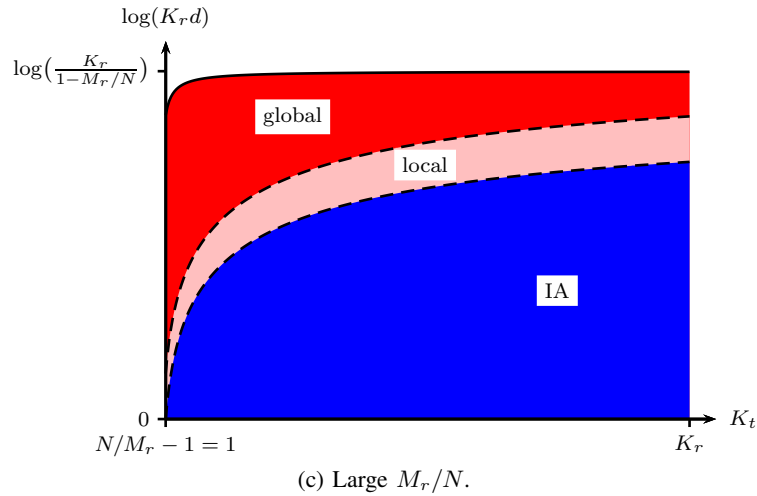
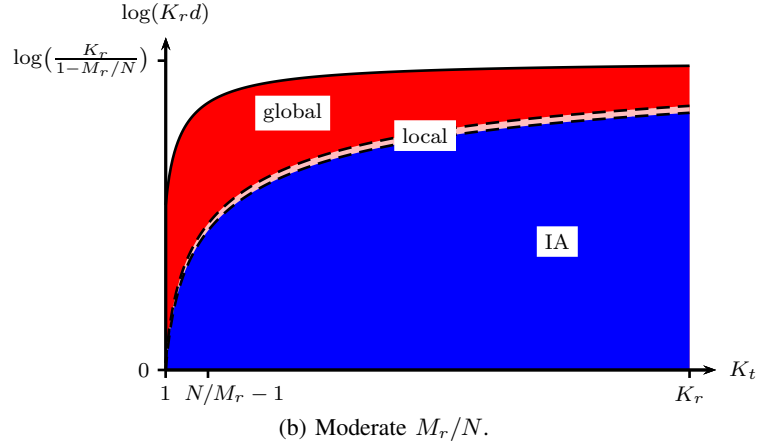
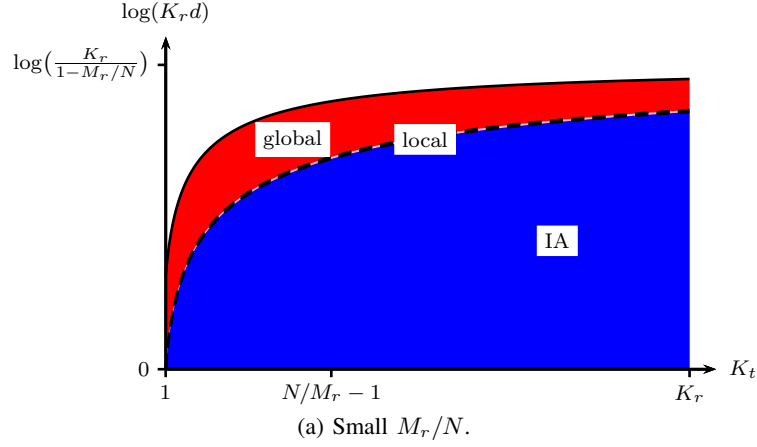


Fig. 6. DoF gains as a function of number of transmitters K_t for various regimes of receiver cache size (characterized by M_r/N).

An interesting question then is how large K_t has to be for the DoF to approach the limit in (6).

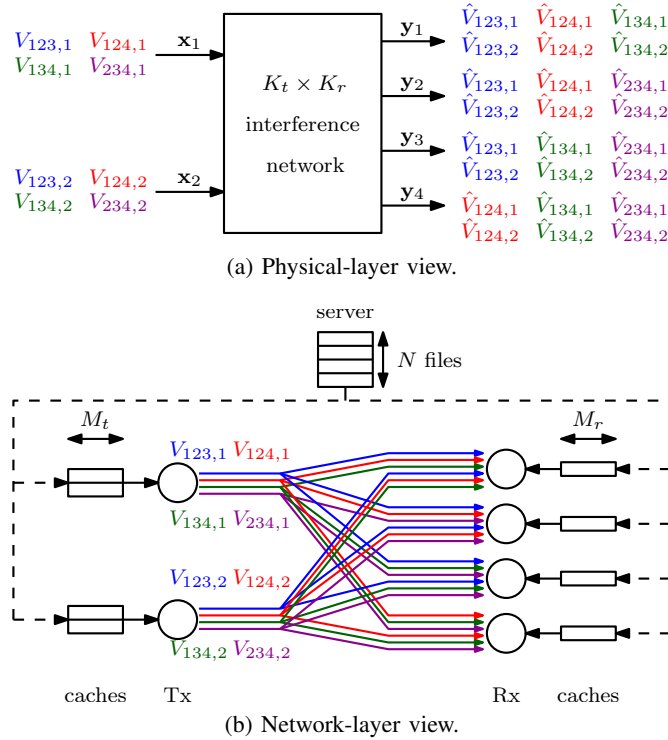


Fig. 7. The separation architecture applied to the setup in Fig. 2 with multicast size $\kappa + 1 = 3$. The interface messages V_{S_j} at the physical layer can be abstracted as orthogonal error-free multicast bit pipes at the network layer. Thus at the physical layer (a) we focus on transmitting the V_{S_j} 's across the interference channel, while at the network layer (b) we perform the caching and delivery strategies, oblivious of the underlying physical channel, to deliver the requested files.

Specifically, for what values of K_t does the sum DoF become $\Theta(\frac{K_r}{1-M_r/N})$? When the receiver cache memory is small, specifically $M_r < N/K_r$, the number of transmitters K_t must be of the order of K_r (see Fig. 6a). However, as M_r increases, we find that a smaller number of transmitters is needed to achieve the same DoF (see Figs. 6b and 6c). In general, the limiting value is reached (within a constant) when $K_t = \Omega(N/M_r - 1)$.² In particular, if the receiver caches can store a constant fraction of the content library, then we only need a constant number of transmitters to achieve maximal benefits, up to a multiplicative constant. There is thus a trade-off between the number of transmitters K_t and the amount of receiver cache memory M_r required for maximal system performance (up to the local caching gain): the larger the receiver memory, the fewer the required transmitters.

While the separation architecture discussed above (on which we focus in most of this paper) is order optimal, one can still make some strict improvements, albeit no more than a constant factor, by choosing a different separation architecture. In Section VII, we present an alternative separation architecture for the case $K_t = K_r = N = 2$ that creates *interacting* error-free bit pipes as the physical-layer abstraction. This architecture can achieve a strictly higher DoF than Theorem 1 in some regimes.

IV. SEPARATION ARCHITECTURE

Our proposed separation architecture isolates the channel coding aspect of the problem from its content delivery aspect. The former is handled at the physical layer, while the latter is handled at the network layer. The two layers interface using a set \mathcal{V} of multiple multicast messages,

$$\mathcal{V} = \{V_{S_j} : j \in \{1, \dots, K_t\}, \mathcal{S} \in \mathcal{S}\}, \quad (7)$$

²This comes from being able to write $K_r \text{DoF} \approx \frac{K_t(\kappa+1)}{K_t(\kappa+1)+K_r-\kappa-1} \cdot \frac{K_r}{1-M_r/N}$, where $\kappa = K_r M_r / N$. The first factor is a constant when $K_t(\kappa+1) = \Omega(K_r - \kappa - 1)$, which leads to $K_t = \Omega(\frac{K_r}{K_r M_r / N + 1} - 1)$. When K_r is large, this behavior becomes $K_t = \Omega(\frac{N}{M_r} - 1)$.

where V_{S_j} denotes the message sent from transmitter j to the subset \mathcal{S} of receivers, and $\mathcal{S} \subseteq 2^{\{1, \dots, K_r\}}$ is some collection of subsets of receivers. Notice that all transmitters have messages for the same subsets of receivers, a natural design choice due to the symmetry of the problem. The physical layer processing transmits these messages across the interference network, while the network layer treats them as orthogonal error-free multicast bit pipes. Fig. 7 illustrates this separation for the setting in Example 1 in Section III.

In order to motivate our choice of \mathcal{S} (and hence of \mathcal{V}), it will be useful to give a brief overview of the strategy used for the broadcast setup in [2]. Suppose that the receiver memory is $M_r = \kappa N / K_r$, where $\kappa \in \{0, 1, \dots, K_r - 1\}$ is an integer. The idea is to place content in the receiver caches such that every subset of κ of them shares an exclusive part of every file (each file is thus split into $\binom{K_r}{\kappa}$ equal parts). During the delivery phase, linear combinations of these file parts are sent to every subset of $\kappa + 1$ users such that each user can combine its received linear combination with the contents of its cache to decode one part of their requested file. As a result, a total of

$$L^{\text{BC}}(N, K_r, M_r) \cdot F = \frac{K_r - \kappa}{\kappa + 1} \cdot F \quad (8)$$

bits are sent through the network (see [2, Theorem 1]).

Notice that the broadcast strategy never really sends any *broadcast* message on a logical level (except when $\kappa + 1 = K_r$). Instead, it sends many *multicast* messages, each intended for $\kappa + 1$ users, which just happen to be “overheard” by the unintended receivers. Inspired by this, we choose the messages in \mathcal{V} to reflect the multicast structure in [2]. Specifically, we choose to create one multicast message from each transmitter to every subset of receivers of size $\kappa + 1$. In other words,

$$\mathcal{S} = \{\mathcal{S} \subseteq \{1, \dots, K_r\} : |\mathcal{S}| = \kappa + 1\}. \quad (9)$$

For example, Fig. 7 shows the separation architecture when $\kappa + 1 = 3$. While (9) depicts the choice of \mathcal{S} that we make most of the time, it is inefficient in a particular regime, namely when both the number of files and the receiver memory are small. Since that regime is of only limited interest, we relegate its description to Appendix A.

Let $\tilde{R}_{\kappa+1}$ be the rate at which we transmit these messages at the physical layer, i.e., $V_{S_j} \in [2^{\tilde{R}_{\kappa+1}T}]$. Further, let $\ell_{\kappa+1}$ be the size (normalized by file size) of whatever is sent through each multicast link at the network layer, i.e., $V_{S_j} \in [2^{\ell_{\kappa+1}F}]$. Therefore, $\tilde{R}_{\kappa+1}T = \ell_{\kappa+1}F$. Let us write $\tilde{R}_{\kappa+1}^*$ and $\ell_{\kappa+1}^*$ to denote the optimal $\tilde{R}_{\kappa+1}$ and $\ell_{\kappa+1}$, respectively, within their respective subproblems (these will be defined rigorously in the subsections below). These quantities can be connected to the rate R of the original caching problem. Indeed, since $F = RT$, then we can achieve a rate R equal to

$$R = \frac{\tilde{R}_{\kappa+1}^*}{\ell_{\kappa+1}^*}, \quad (10)$$

when $M_r = \kappa N / K_r$, $\kappa \in \{0, 1, \dots, K_r - 1\}$.³

The separation architecture has thus created two subproblems of the original problem. At the physical layer, we have a pure communication subproblem, where multicast messages V_{S_j} must be transmitted reliably across an interference network. At the network layer, we have a caching subproblem with noiseless orthogonal multicast links connecting transmitters to receivers. In the two subsections below, we properly formulate each subproblem. We give a strategy for each as well as the values of $\tilde{R}_{\kappa+1}$ and $\ell_{\kappa+1}$ that they achieve.

³The nature of the separation architecture implies that κ must always be an integer. Regimes where it is not are handled using time and memory sharing between points where it is. Furthermore, we exclude the case $\kappa = K_r$ (equivalently, $M_r = N$) for mathematical convenience, but we can in fact trivially achieve an infinite rate when $M_r = N$ by storing the complete content library in every user’s cache.

A. Physical Layer

At the physical layer, we consider only the communication problem of transmitting specific messages across the interference channel described in Section II, as illustrated in Fig. 7a. This is an interesting communication problem on its own, and we hence formulate it without all the caching details. The message set that we consider is one where every transmitter j has a message for every subset \mathcal{S} of σ receivers, where $\sigma \in \{1, \dots, K_r\}$ is given.⁴ We label such a message as $V_{\mathcal{S}j}$, and we note that there are a total of $K_t \binom{K_r}{\sigma}$ of them. For instance, in the example shown in Fig. 7a, message $V_{134,2}$ (used as a shorthand for $V_{\{1,3,4\},2}$) is sent by transmitter 2 to receivers 1, 3, and 4. We call this problem the multiple multicast X-channel with multicast size σ , as it generalizes the (unicast) X-channel studied in [4] to multicast messages. Note that, when $\sigma = 1$, we recover the unicast X-channel.

We assume a symmetric setup, where all the messages have the same rate \tilde{R}_σ , i.e., $V_{\mathcal{S}j} \in [2^{\tilde{R}_\sigma T}]$. A rate is called achievable if a strategy exists allowing all receivers to recover all their intended messages with vanishing error probability as the block length T increases. Our goal is to find the largest achievable rate \tilde{R}_σ for a given SNR, denoted by $\tilde{R}_\sigma^*(\text{SNR})$, and in particular its DoF

$$\tilde{d}_\sigma^*(K_t, K_r) \triangleq \lim_{\text{SNR} \rightarrow \infty} \frac{\tilde{R}_\sigma^*(\text{SNR})}{\frac{1}{2} \log \text{SNR}}.$$

One of the contributions of this paper is an exact characterization of \tilde{d}_σ^* , and we next give an overview of how to achieve it.

For every receiver i , there is a set of $K_t \binom{K_r-1}{\sigma-1}$ desired messages $\{V_{\mathcal{S}j} : i \in \mathcal{S}\}$, and a set of $K_t \binom{K_r-1}{\sigma}$ interfering messages $\{V_{\mathcal{S}j} : i \notin \mathcal{S}\}$. Using TDMA, all $K_t \binom{K_r}{\sigma}$ messages can be delivered to their receivers at a sum DoF of 1, i.e., $\tilde{d}_\sigma = \frac{1}{K_t \binom{K_r}{\sigma}}$. However, by applying an interference alignment technique that generalizes the one used in [4], we can, loosely speaking, collapse the $K_t \binom{K_r-1}{\sigma}$ interfering messages at every receiver into a subspace of dimension $\binom{K_r-1}{\sigma}$ (assuming for simplicity that each message forms a subspace of dimension one), while still allowing reliable recovery of all $K_t \binom{K_r-1}{\sigma-1}$ desired messages. Thus an overall vector space of dimension $K_t \binom{K_r-1}{\sigma-1} + \binom{K_r-1}{\sigma} < K_t \binom{K_r}{\sigma}$ is used to deliver all $K_t \binom{K_r}{\sigma}$ messages. This strategy achieves a DoF-optimal rate, as asserted by the following theorem.

Theorem 2. *The DoF of the symmetric multiple multicast X-channel with multicast size σ is given by*

$$\tilde{d}_\sigma^*(K_t, K_r) = \frac{1}{K_t \binom{K_r-1}{\sigma-1} + \binom{K_r-1}{\sigma}}.$$

The details of the interference alignment strategy are given in Section V. The proof of optimality is left for Appendix D, since it does not directly contribute to our main result in Theorem 1. It does however reinforce it by providing a complete solution to the physical-layer communication subproblem.

The DoF shown in Theorem 2 is a per-message DoF. Since there are a total of $K_t \binom{K_r}{\sigma}$ messages, we obtain a sum DoF of

$$K_t \binom{K_r}{\sigma} \cdot \tilde{d}_\sigma^*(K_t, K_r) = \frac{K_t \binom{K_r}{\sigma}}{K_t \binom{K_r-1}{\sigma-1} + \binom{K_r-1}{\sigma}} = \frac{K_t K_r}{(K_t - 1)\sigma + K_r}. \quad (11)$$

When $\sigma = 1$, the sum DoF in (11) is $\frac{K_t K_r}{K_t + K_r - 1}$, thus recovering the unicast X-channel result from [4]. When $\sigma = K_r$, the problem reduces to a broadcast channel with multiple sources, giving a sum DoF of 1.

⁴In the context of the caching problem, σ is chosen to be $\kappa + 1$, as described earlier.

B. Network Layer

The network layer setup is similar to the end-to-end setup, with the difference that the interference network is replaced by the multicast links V_{S_j} from transmitters to receivers, as illustrated in Fig. 7b. As mentioned previously, each link V_{S_j} is shared by exactly $|\mathcal{S}| = \kappa + 1$ users, where $\kappa = K_r M_r / N$ is an integer. We again focus on a symmetric setup, where all links have the same size $\ell_{\kappa+1}$, called the *link load*. It will be easier in the discussion to use the *sum network load* $L_{\kappa+1}$, i.e., the combined load of all $K_t \binom{K_r}{\kappa+1}$ links,

$$L_{\kappa+1} = K_t \binom{K_r}{\kappa+1} \cdot \ell_{\kappa+1}. \quad (12)$$

A sum network load L is said to be achievable if, for every large enough file size F , a strategy exists allowing all users to recover their requested files with high probability while transmitting no more than LF bits through the network. Our goal is to find the smallest achievable network load for every N , K_t , K_r , M_t , and M_r , denoted by

$$L_{\kappa+1}^*(N, K_t, K_r, M_t, M_r),$$

where $\kappa = K_r M_r / N$ is an integer. Using a similar strategy to [2], we achieve the following sum network load.

Lemma 3. *In the network layer setup with a multicast size of $\kappa + 1$, $\kappa \in \{0, 1, \dots, K_r - 1\}$, a sum network load of*

$$L_{\kappa+1}^*(N, K_t, K_r, M_t, M_r) \leq \frac{K_r - \kappa}{\kappa + 1}$$

can be achieved when $M_r = \kappa N / K_r$.

Proof: We first divide every file W_n into K_t equal parts, $W_n = (W_n^1, \dots, W_n^{K_t})$, and store the j -th part W_n^j in the cache of transmitter j . Note that, while we allow $M_t \geq N / K_t$ as per the regularity condition in (1), the above transmitter placement only stores exactly N / K_t files at every transmitter irrespective of the value of M_t . The different transmitters are then treated as independent sublibraries. Indeed, the receiver placement splits each receiver cache into K_t equal sections, and each section is dedicated to one sublibrary. A placement phase identical to [2] is then performed for each sublibrary in its dedicated receiver memory.

During the delivery phase, user i 's request for a single file W_{u_i} is converted into K_t separate requests of the subfiles $(W_{u_i}^1, \dots, W_{u_i}^{K_t})$, each from its corresponding sublibrary (transmitter). For every subset \mathcal{S} of $\kappa + 1$ receivers, each transmitter j then sends through the link V_{S_j} exactly what would be sent to these receivers in the broadcast setup, had the other transmitters not existed. This is possible since the V_{S_j} links were chosen by design to match the multicast transmissions in the broadcast setup. Each transmitter will thus send $\frac{1}{K_t} \cdot L^{\text{BC}}(N, K_r, M_r)$ files through the network (with L^{BC} as defined in (8)), for a total network load of $\frac{K_r - \kappa}{\kappa + 1}$. ■

C. Achievable End-to-End DoF

From (10) and using (12), we can achieve an end-to-end DoF of

$$\text{DoF} \geq \frac{\tilde{d}_{\kappa+1}^*(K_t, K_r)}{L_{\kappa+1}^*(N, K_t, K_r, M_t, M_r)} \cdot K_t \binom{K_r}{\kappa+1}.$$

By combining Theorem 2 (with $\sigma = \kappa + 1$) and Lemma 3,

$$\begin{aligned} \frac{1}{\text{DoF}} &\leq \frac{L_{\kappa+1}^*(N, K_t, K_r, M_t, M_r)}{\tilde{d}_{\kappa+1}^*(K_t, K_r)} \cdot \frac{1}{K_t \binom{K_r}{\kappa+1}} \\ &\leq \frac{K_r - \kappa}{\kappa + 1} \cdot \left[K_t \binom{K_r - 1}{\kappa} + \binom{K_r - 1}{\kappa + 1} \right] \cdot \frac{1}{K_t \binom{K_r}{\kappa+1}}. \end{aligned}$$

By writing

$$\begin{aligned}
K_t \binom{K_r - 1}{\kappa} + \binom{K_r - 1}{\kappa + 1} &= (K_t - 1) \binom{K_r - 1}{\kappa} + \binom{K_r - 1}{\kappa} + \binom{K_r - 1}{\kappa + 1} \\
&\stackrel{(a)}{=} (K_t - 1) \frac{\kappa + 1}{K_r} \binom{K_r}{\kappa + 1} + \binom{K_r}{\kappa + 1} \\
&= \frac{(K_t - 1)(\kappa + 1) + K_r}{K_r} \cdot \binom{K_r}{\kappa + 1},
\end{aligned}$$

where (a) is due to Pascal's triangle, we conclude that

$$\frac{1}{\text{DoF}} \leq \frac{K_r - \kappa}{\kappa + 1} \cdot \frac{(K_t - 1)(\kappa + 1) + K_r}{K_t K_r} = \frac{K_t - 1 + \frac{K_r}{\kappa + 1}}{K_t} \cdot \left(1 - \frac{\kappa}{K_r}\right). \quad (13)$$

This proves the achievability direction of Theorem 1 when $\frac{K_r}{\kappa + 1} \leq N$. The case $\frac{K_r}{\kappa + 1} > N$ is discussed in Appendix A.

V. THE MULTIPLE MULTICAST X-CHANNEL

The multiple multicast X-channel problem (with multicast size σ) that emerges from our separation strategy is a generalization of the unicast ($\sigma = 1$) X-channel studied in [4]. We propose an interference alignment strategy that generalizes the one in [4]. In this section, we give a high-level overview of the alignment strategy in order to focus on the intuition. The rigorous explanation of the strategy is given in Appendix B as a proof of Lemma 4, which is presented at the end of this section.

Consider communicating across the interference network over T time slots. Every transmitter j beamforms each message V_{S_j} along some fixed vector of length T and sends the sum of the vectors corresponding to all its messages as its codeword. Each message thus occupies a subspace of dimension 1 of the overall T -dimensional vector space. The goal is to align at each receiver the interfering messages into the smallest possible subspace, so that a high rate is achieved for the desired messages.

When choosing which messages to align, we enforce the following three principles, which ensure maximal alignment without preventing decodability of the intended messages. At every receiver i :

- 1) Each desired message V_{S_j} with $i \in \mathcal{S}$ must be in a subspace of dimension 1, not aligned with any other subspace.
- 2) Messages from the same transmitter must never be aligned.
- 3) All messages intended for the same subset \mathcal{S} of receivers with $i \notin \mathcal{S}$ must be aligned into one subspace of dimension 1.

Principle 1 ensures that receiver i can decode all of its desired messages. To understand principle 2, notice that messages from the same transmitter go through the same channels. Therefore, if two messages from the same transmitter are aligned at one receiver, then they were also aligned during transmission, and are hence aligned at all other receivers, including their intended ones. Thus principle 2 ensures decodability at other receivers. As for principle 3, it provides the maximal alignment of the interfering messages without violating principle 2. Indeed, each aligned subspace contains K_t messages, one from each transmitter. Any additional message that is aligned would share a transmitter with one of them.

For every receiver, there are $K_t \binom{K_r - 1}{\sigma - 1}$ desired messages. By principle 1, each should take up one non-aligned subspace of dimension 1, for a total of $K_t \binom{K_r - 1}{\sigma - 1}$ dimensions. On the other hand, there are $K_t \binom{K_r - 1}{\sigma}$ interfering messages. By principle 3, every K_t of them are aligned in one subspace of dimension 1, and hence all interfering messages fall in a subspace of dimension $\binom{K_r - 1}{\sigma}$. These subspaces can be made non-aligned by ensuring that the overall vector space has a dimension of

$$T = K_t \binom{K_r - 1}{\sigma - 1} + \binom{K_r - 1}{\sigma}.$$

Since each message took up one dimension, we get a per-message DoF of

$$\frac{1}{T} = \frac{1}{K_t \binom{K_r-1}{\sigma-1} + \binom{K_r-1}{\sigma}}.$$

This is an improvement over TDMA, which achieves a DoF of $1/K_t \binom{K_r}{\sigma}$.

In most cases, we do not achieve the exact DoF shown in Theorem 2 using a finite number of channel realizations. We instead achieve an arbitrarily close DoF by using an increasing number of channel realizations. The exact achieved DoF is given in the following lemma.

Lemma 4. *Let $\Gamma = (K_r - \sigma)(K_t - 1)$. For any arbitrary $n \in \mathbb{N}$, we can achieve a DoF for message V_{Sj} equal to*

$$\delta_j^{(n)} = \frac{(n + c_j)^\Gamma}{\binom{K_r-1}{\sigma-1} [(n+1)^\Gamma + (K_t-1)n^\Gamma] + \binom{K_r-1}{\sigma} (n+1)^\Gamma},$$

where $c_1 = 1$ and $c_2 = \dots = c_{K_t} = 0$.

The proof of Lemma 4 is given in Appendix B.

Note that Lemma 4 achieves a slightly different DoF for V_{Sj} depending on j , which might seem to contradict the symmetry in the problem setting. However, for a large n , we have $(n+1)^\Gamma \approx n^\Gamma$, and hence

$$\lim_{n \rightarrow \infty} \delta_j^{(n)} = \frac{1}{K_t \binom{K_r-1}{\sigma-1} + \binom{K_r-1}{\sigma}} = \tilde{d}_\sigma^*(K_t, K_r)$$

for all j . Thus the symmetric DoF $\tilde{d}_\sigma^*(K_t, K_r)$ is achieved in the limit.

VI. ORDER-OPTIMALITY OF THE SEPARATION ARCHITECTURE

In this section, we give a high-level proof of the converse part of Theorem 1 by showing that the DoF achieved by the separation architecture in Section IV is order-optimal. We do this by computing cut-set-based information-theoretic upper bounds on the DoF (equivalently, they are lower bounds on the reciprocal $1/\text{DoF}$). These bounds are given in the following lemma, whose proof is placed at the end of this section in order not to distract from the intuition behind the converse arguments. The rigorous converse proof is given in Appendix C.

Lemma 5. *For any $N, K_t, K_r, M_t \in [0, N]$, and $M_r \in [0, N]$, the optimal DoF must satisfy*

$$\frac{1}{\text{DoF}} \geq \max_{s \in \{1, \dots, \min\{K_r, N\}\}} \frac{s \left(1 - \frac{M_r}{\lfloor N/s \rfloor}\right)}{\min\{s, K_t\}}.$$

Lemma 5 is next used to prove the converse part of Theorem 1, i.e.,

$$\text{DoF} \leq 13.5d(N, K_t, K_r, M_t, M_r),$$

where $d(\cdot)$ is defined in (4). The procedure is similar to the one used in [2]: we consider three main regimes (Regimes 1, 2, and 3) of receiver memory M_r and in each compare the expression $d(\cdot)$ with the outer bounds. In addition, we consider a separate corner case (Regime 0) in which the largest possible number of distinct file requests (i.e., $\min\{K_r, N\}$) is small compared to the number of transmitters.

$$\text{Regime 0: } \min\{K_r, N\} \leq 12.5K_t; \tag{14a}$$

$$\text{Regime 1: } \min\{K_r, N\} > 12.5K_t \quad \text{and} \quad 0 \leq M_r \leq 1.1 \max \left\{ 1, \frac{N}{K_r} \right\}; \tag{14b}$$

$$\text{Regime 2: } \min\{K_r, N\} > 12.5K_t \quad \text{and} \quad 1.1 \max \left\{ 1, \frac{N}{K_r} \right\} < M_r \leq 0.092 \frac{N}{K_t}; \tag{14c}$$

$$\text{Regime 3: } \min\{K_r, N\} > 12.5K_t \quad \text{and} \quad 0.092 \frac{N}{K_t} < M_r \leq N. \tag{14d}$$

Note that Regimes 1, 2, and 3 are unambiguous, since

$$\min\{K_r, N\} > 12.5K_t \implies 0 < 1.1 \max\left\{1, \frac{N}{K_r}\right\} < 0.092 \frac{N}{K_t} < N. \quad (15)$$

Since M_r is the only variable that we will consistently vary, we will abuse notation for convenience and write $d(M_r)$ instead of $d(N, K_t, K_r, M_t, M_r)$ for all $M_r \in [0, N]$. Our goal is thus to prove

$$\frac{1}{\text{DoF}} \geq \frac{d^{-1}(M_r)}{13.5}. \quad (16)$$

For ease of reference, we will rewrite the expression of $d^{-1}(M_r)$ here. For $M_r = \kappa N/K_r$ where $\kappa \in \{0, 1, \dots, K_r\}$ is an integer,

$$d^{-1}(\kappa N/K_r) = \frac{K_t - 1 + \min\left\{\frac{K_r}{\kappa+1}, N\right\}}{K_t} \cdot \left(1 - \frac{\kappa}{K_r}\right), \quad (17)$$

and $d^{-1}(M_r)$ is the lower convex envelope of these points for all $M_r \in [0, N]$. Note that $d^{-1}(M_r)$ is non-increasing and convex in M_r .

Regimes 0 and 3: Interestingly, Regimes 0 and 3 behave quite similarly to each other. Indeed, notice that in both of them we have $K_t = \Omega\left(\min\left\{\frac{K_r}{\kappa+1}, N\right\}\right)$. Using (17), this implies

$$d^{-1}(M_r) \approx \frac{K_t + \min\left\{\frac{K_r}{\kappa+1}, N\right\}}{K_t} \cdot \left(1 - \frac{M_r}{N}\right) = \Theta\left(1 - \frac{M_r}{N}\right).$$

Conversely, we can apply Lemma 5 with $s = 1$ to get $1/\text{DoF} \geq 1 - M_r/N$. Thus in both regimes the local caching gain is the only significant contribution to the DoF.

Regime 1: In Regime 1, the receiver memory is too small to have any significant effect. Therefore, using $12.5K_t < \min\{K_r, N\}$, we can write (17) as

$$d^{-1}(M_r) \approx \frac{K_t + \min\left\{\frac{K_r}{\kappa+1}, N\right\}}{K_t} \cdot 1 \leq \left(\frac{1}{12.5} + 1\right) \cdot \frac{\min\{K_r, N\}}{K_t}.$$

Conversely, by using Lemma 5 with $s \approx \min\{K_r, N\}$, we get

$$\frac{1}{\text{DoF}} \gtrsim \frac{s \cdot 1}{K_t} \approx \frac{\min\{K_r, N\}}{K_t}.$$

Therefore, in this regime $\text{DoF} \approx \frac{K_t}{\min\{K_r, N\}}$. We can explain this in terms of the DoF gains in (5): when the receiver memory is very small, the only relevant gain is the interference alignment gain.

Regime 2: In Regime 2, the receivers combined can store all of the content library. As a result, the global caching gain kicks in. We can upper-bound $d^{-1}(M_r)$ in (17) as follows:

$$d^{-1}(M_r) \leq 1 + \frac{K_r}{K_t(K_r M_r/N + 1)} \leq 1 + \frac{N}{K_t M_r} \leq 1.092 \frac{N}{K_t M_r},$$

because $K_t < 0.092N/M_r$ in Regime 2. Conversely, let us apply Lemma 5 using $s \approx \frac{N}{2M_r}$:

$$\frac{1}{\text{DoF}} \gtrsim \frac{s - s^2 M_r/N}{K_t} \approx \frac{N}{4K_t M_r}.$$

Therefore, $\text{DoF} \approx \frac{K_t M_r}{N}$. This behavior is similar to what one would expect in the broadcast setup in [2], with the exception of the additional K_t factor.

Since $d^{-1}(M_r)$ approximately matches the outer bounds in all four regimes and can also be achieved as in Section IV, then it provides an approximate characterization of $1/\text{DoF}$. The above arguments are made rigorous in Appendix C.

Proof of Lemma 5: Consider $s \in \{1, \dots, \min\{K_r, N\}\}$ users. We shall look at $E = \lfloor N/s \rfloor$ different request vectors, such that the combined number of files requested by all users after E request instances is $\tilde{N} = sE = s \lfloor N/s \rfloor$ files. More specifically, we consider the request vectors $\mathbf{u}_1, \dots, \mathbf{u}_E$ with

$$\mathbf{u}_e = \left(\underbrace{(e-1)s+1, (e-1)s+2, \dots, es}_s, \underbrace{1, \dots, 1}_{K_r-s} \right),$$

for each $e = 1, \dots, E$. Note that we only focus on the first s users; the remaining $K_r - s$ users are not relevant to our argument.

When the request vector is \mathbf{u} , let $\mathbf{x}_j^{\mathbf{u}}$ and $\mathbf{y}_i^{\mathbf{u}}$ denote the inputs and outputs of the interference network for all transmitters j and receivers i . For notational convenience, we write $\mathbf{y}_{[s]}^{\mathbf{u}} = (\mathbf{y}_1^{\mathbf{u}}, \dots, \mathbf{y}_s^{\mathbf{u}})$ and use a similar notation for $\mathbf{x}_{[K_t]}^{\mathbf{u}}$. Also, let Q_i denote the contents of user i 's cache (recall that the cache contents are independent of \mathbf{u}). By Fano's inequality,

$$H\left(W_1, \dots, W_{\tilde{N}} \middle| Q_1, \dots, Q_s, \mathbf{y}_{[s]}^{\mathbf{u}_1}, \dots, \mathbf{y}_{[s]}^{\mathbf{u}_E}\right) \leq \varepsilon T, \quad (18)$$

since the s users should be able to each decode their $\lfloor N/s \rfloor$ requested files using their caches and channel outputs. Then,

$$\begin{aligned} \tilde{N}RT &= H(W_1, \dots, W_{\tilde{N}}) \\ &= I\left(W_1, \dots, W_{\tilde{N}}; Q_1, \dots, Q_s, \mathbf{y}_{[s]}^{\mathbf{u}_1}, \dots, \mathbf{y}_{[s]}^{\mathbf{u}_E}\right) + H\left(W_1, \dots, W_{\tilde{N}} \middle| Q_1, \dots, Q_s, \mathbf{y}_{[s]}^{\mathbf{u}_1}, \dots, \mathbf{y}_{[s]}^{\mathbf{u}_E}\right) \\ &\stackrel{(a)}{\leq} I\left(W_1, \dots, W_{\tilde{N}}; Q_1, \dots, Q_s, \mathbf{y}_{[s]}^{\mathbf{u}_1}, \dots, \mathbf{y}_{[s]}^{\mathbf{u}_E}\right) + \varepsilon T \\ &= I\left(W_1, \dots, W_{\tilde{N}}; \mathbf{y}_{[s]}^{\mathbf{u}_1}, \dots, \mathbf{y}_{[s]}^{\mathbf{u}_E}\right) + I\left(W_1, \dots, W_{\tilde{N}}; Q_1, \dots, Q_s \middle| \mathbf{y}_{[s]}^{\mathbf{u}_1}, \dots, \mathbf{y}_{[s]}^{\mathbf{u}_E}\right) + \varepsilon T \\ &\leq I\left(W_1, \dots, W_{\tilde{N}}; \mathbf{y}_{[s]}^{\mathbf{u}_1}, \dots, \mathbf{y}_{[s]}^{\mathbf{u}_E}\right) + H(Q_1, \dots, Q_s) + \varepsilon T \\ &\stackrel{(b)}{\leq} I\left(\mathbf{x}_{[K_t]}^{\mathbf{u}_1}, \dots, \mathbf{x}_{[K_t]}^{\mathbf{u}_E}; \mathbf{y}_{[s]}^{\mathbf{u}_1}, \dots, \mathbf{y}_{[s]}^{\mathbf{u}_E}\right) + sM_rRT + \varepsilon T \\ &\stackrel{(c)}{\leq} E \cdot \max_{e \in \{1, \dots, E\}} I\left(\mathbf{x}_1^{\mathbf{u}_e}, \dots, \mathbf{x}_{K_t}^{\mathbf{u}_e}; \mathbf{y}_1^{\mathbf{u}_e}, \dots, \mathbf{y}_s^{\mathbf{u}_e}\right) + sM_rRT + \varepsilon T \\ &\stackrel{(d)}{\leq} E \cdot T \left(\min\{K_t, s\} \cdot \frac{1}{2} \log \text{SNR} + o(\log \text{SNR}) \right) + sM_rRT + \varepsilon T \\ &= \lfloor N/s \rfloor T \cdot \left(\min\{K_t, s\} \cdot \frac{1}{2} \log \text{SNR} + o(\log \text{SNR}) \right) + sM_rRT + \varepsilon T, \end{aligned}$$

where (a) is due to inequality (18), (b) uses the data processing inequality, (c) follows from the independence of the channel outputs when conditioned on all channel inputs, and (d) is the capacity bound of the $K_t \times s$ MIMO channel over T time blocks.

Since $\tilde{N} = s \lfloor N/s \rfloor$, and by taking $T \rightarrow \infty$ and $\varepsilon \rightarrow 0$, we obtain

$$R \left(1 - \frac{M_r}{\lfloor N/s \rfloor} \right) \leq \frac{1}{s} \cdot \min\{K_t, s\} \cdot \frac{1}{2} \log \text{SNR} + o(\log \text{SNR}).$$

The DoF thus obeys

$$\text{DoF} \leq \frac{\min\{K_t, s\}}{s \left(1 - \frac{M_r}{\lfloor N/s \rfloor} \right)}.$$

Since s was arbitrary, the above is true for any $s \in \{1, \dots, \min\{K_r, N\}\}$, and thus the lemma is proved. \blacksquare

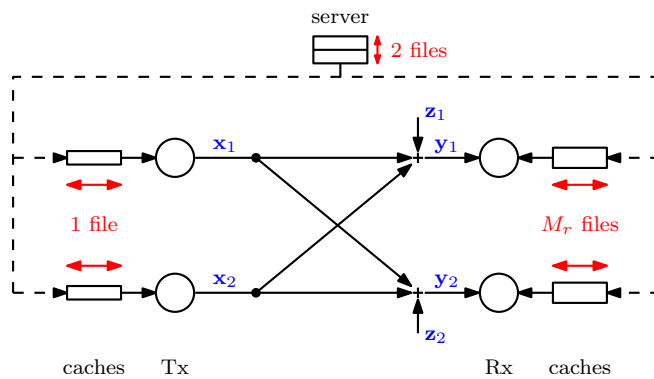


Fig. 8. The 2×2 cache-aided interference channel with 2 files. The transmitter caches can hold exactly one file each, and the receiver caches $M_r \in [0, 2]$ files each. The \mathbf{z}_i 's are iid additive Gaussian unit-variance noise.

VII. DISCUSSION AND EXTENSIONS

In this paper, we have determined the approximate DoF of the general cache-aided interference network. To do so, we have proposed a separation-based strategy that uses interference alignment to create *non-interacting* multicast bit pipes from transmitters to receivers, and we have shown that this strategy achieves a DoF that is within a constant multiplicative factor from the optimum. However, this achieved DoF is only approximately optimal. In fact, many improvements can be made, such as using transmit zero-forcing as has been discussed in previous work [3], [34], [35]. In this section, we explore a different approach, which lies within the context of interference alignment described in Section V: rather than ignoring the interference subspace, which contains the aligned messages, we attempt to extract some information from it. Thus every receiver gains additional information in the form of an alignment of the bit pipes available at other receivers: the bit pipes would thus *interact*. We study this approach in a specific setup: the 2×2 interference channel with a content library of two files, shown in Fig. 8.

For this 2×2 setup, by Theorem 1 the main strategy described in this paper achieves

$$\frac{1}{\text{DoF}} \leq \max \left\{ \frac{3}{2} - M_r, 1 - \frac{1}{2}M_r \right\},$$

for $M_r \in [0, 2]$, as shown by the solid line in Fig. 9. However, the same figure shows an improved inverse DoF, depicted by the dashed line, which is achieved using the interference-extracting scheme discussed in this section. A factor-7/6 improvement is obtained over the main strategy. This result is summarized in the following theorem.

Theorem 6. *The following inverse DoF can be achieved for the 2×2 cache-aided interference network with $N = 2$ files and transmitter memory $M_t = 1$:*

$$\frac{1}{\text{DoF}} \leq \max \left\{ \frac{3}{2} - \frac{3}{2}M_r, \frac{9}{7} - \frac{6}{7}M_r, 1 - \frac{1}{2}M_r \right\},$$

for all values of $M_r \in [0, 2]$.

It should be noted that the general converse stated in Lemma 5 can be applied here and results in

$$\frac{1}{\text{DoF}} \geq 1 - \frac{1}{2}M_r,$$

which implies that our strategy is exactly optimal for $M_r \geq 4/5$, as illustrated by the dash-dotted line in Fig. 9.

We will next give a high-level overview of the interference-extraction strategy. The proof of Theorem 6, including the details of the strategy, are given in Appendix F. Consider what happens when the main strategy is used in this 2×2 setup with $M_r = 0$. The strategy creates one unicast message from every

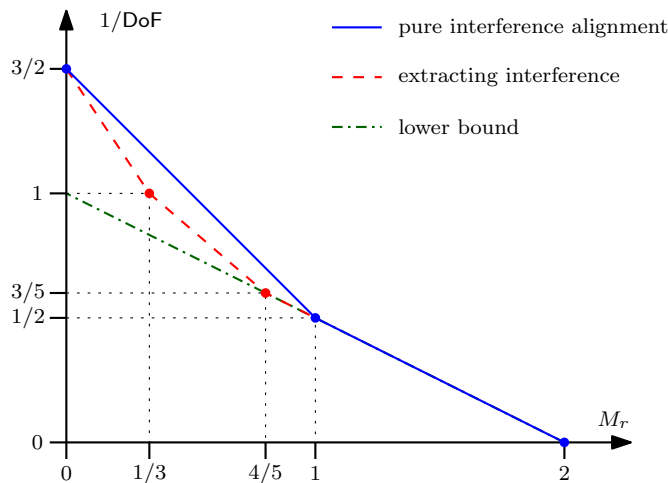


Fig. 9. Inverse DoF achieved by the scheme from Section V (solid line), and the improved inverse DoF achieved by extracting more information from the aligned interference (dashed line). The dash-dotted line shows the information-theoretic lower bounds from Lemma 5.

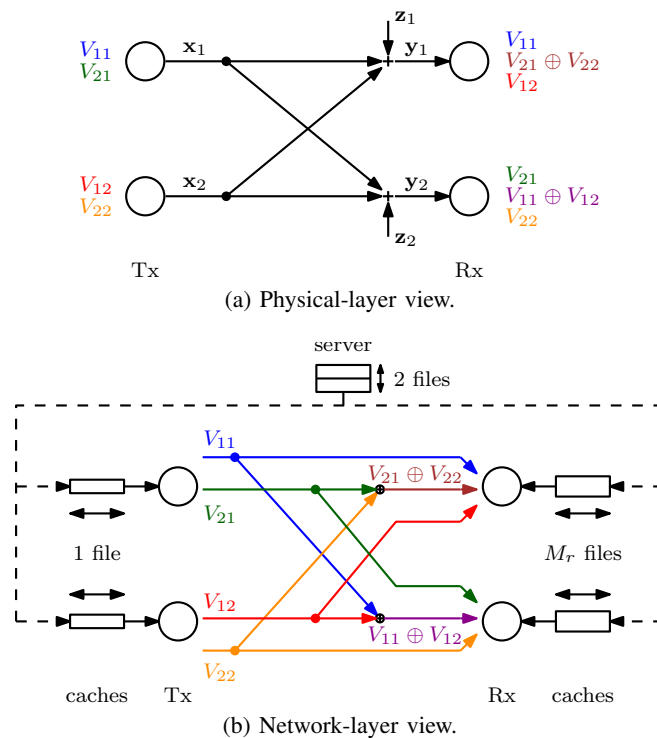


Fig. 10. Separation architecture with interference extraction in the 2×2 case with 2 files. The (unicast) X-channel message set is used, but every receiver decodes, in addition to its intended messages, the sum of the messages intended for the other receiver.

transmitter to every receiver, and transmits them using interference alignment. Each receiver thus gets the two messages intended for it, plus an alignment of the two messages intended for the other receiver. In the main strategy, this aligned interference is simply discarded. However, we can design the scheme in a way that this alignment is a simple sum of the two interfering messages. Each receiver can then decode, in addition to its intended messages, the sum of the interfering messages, without suffering any decrease in the sum DoF of the communicated messages. We hence obtain a new separation architecture, illustrated in Fig. 10, that we use for all M_r .

The scheme we propose in this section is very specific to the 2×2 interference network with two files in the content library. An interesting direction for future work would be to extend this interference-extraction strategy to more general settings.

APPENDIX A
SPECIAL CASE: SMALL NUMBER OF FILES

Recall that the separation architecture creates a set of messages \mathcal{V} as an interface between the physical and network layers,

$$\mathcal{V} = \{V_{Sj} : j \in \{1, \dots, K_t\}, \mathcal{S} \in \mathcal{S}\}$$

for some $\mathcal{S} \subseteq 2^{\{1, \dots, K_r\}}$, as seen in (7). In this paper, we have so far focused on the choice of messages described by \mathcal{S} in (9), in which every transmitter has a message for every subset of exactly $\kappa + 1$ receivers, where $\kappa = K_r M_r / N$ is an integer. While this is order-optimal in most cases, it is insufficient when both the receiver memory and the number of files is small.

To illustrate, consider the case with only a single file in the content library ($N = 1$) and without receiver caches ($M_r = 0$). Furthermore, assume that there is just one transmitter ($K_t = 1$) but many receivers (K_r is large). Seeing as there is only one file, all receivers will request that same file, and hence the obvious strategy is for the transmitter to broadcast the file to all receivers, thus achieving a DoF of 1. However, under the separation architecture described by (9), we create one message from the transmitter for every individual receiver, and then send that file separately as K_r different messages. This is clearly inefficient since the same file is being sent K_r times, thus achieving a much worse DoF of $1/K_r$.

The reason the usual separation architecture is inefficient in this example is that it inherently assumes that all users request different files in the worst case. This is true when there are more files than users. However, if there are so few files that many users will inevitably request the same file, then the previous assumption fails. In this appendix, we handle that case by providing a different separation interface. We exclusively work with the case $M_r = 0$ and compute an achievable DoF for it. Specifically, we show that

$$\frac{1}{\text{DoF}} \leq \frac{K_t + \min\{K_r, N\} - 1}{K_t}. \quad (19)$$

Since we also know that $1/\text{DoF}$, which is convex in M_r , is zero when $M_r = N$, then we can achieve any linear combination of the two reciprocal DoFs between these two points, using time- and memory-sharing. Specifically, we achieve

$$\frac{1}{\text{DoF}} \leq \frac{K_t + \min\{K_r, N\} - 1}{K_t} \cdot \left(1 - \frac{M_r}{N}\right). \quad (20)$$

The expression of the reciprocal DoF in (20) can be decomposed into two gains, in a similar way as in (5). Since the strategy that achieves (20) is relevant when $N < K_r$, we can write the two gains as

$$N\text{DoF} \approx \underbrace{\frac{K_t N}{K_t + N - 1}}_{g^{\text{IA}}} \cdot \underbrace{\frac{1}{1 - \frac{M_r}{N}}}_{g^{\text{LC}}}.$$

Note that $N\text{DoF}$ is the sum DoF here since the total number of requested files is $N < K_r$ in the worst case.

The most striking difference with (5) is that there is no global caching gain. Indeed, the strategy makes no use of any coding or multicasting opportunities, as we will see below. On the other hand, the local caching gain is present and is the same as before. The interference alignment gain is slightly different: it is the interference alignment gain of a $K_t \times N$ unicast X-channel, not $K_t \times K_r$. The reason for this is that, when $N < K_r$, then the total number of *distinct* requested files is N in the worst case. The strategy thus only needs to account for N distinct demands, and uses methods from the compound X-channel [37], [38] to serve them.

We proceed with the strategy for $M_r = 0$ that achieves (19). Since $M_r = 0$, we cannot store anything in the receiver caches. In the transmitter caches, we place the same content as previously described, i.e., every file W_n is split into K_t parts, and transmitter j stores the j -th part of W_n , called W_n^j , for every n . In the delivery phase, we partition the set of users into subsets such that all the users in the same subset

request the same file. Specifically, let \mathbf{u} denote the request vector, and let $\mathcal{U}_n = \{i : u_i = n\}$ be the set of users requesting file W_n . Our goal is to create a multicast message from every transmitter to all users that are requesting the same file. In other words, we set

$$\mathcal{S} = \{\mathcal{U}_n : n \in \{1, \dots, N\} \text{ s.t. } \mathcal{U}_n \neq \emptyset\}$$

Note that \mathcal{S} is a partition of the entire set of users. We denote its size by $\tilde{N} = |\mathcal{S}|$, which is equivalent to the total number of distinct requested files. Our separation interface \mathcal{V} is thus a set of messages from every transmitter to \tilde{N} *non-overlapping* subsets of receivers,

$$\mathcal{V} = \{V_{\mathcal{U}_n j} : \mathcal{U}_n \neq \emptyset \text{ and } j = 1, \dots, K_t\}.$$

We focus on transmitting these messages across the interference channel at the physical layer. At the network layer, we use these messages as error-free bit pipes to deliver the requested files to the users at the network layer.

A. Physical Layer

At the physical layer, the problem is equivalent to the compound X-channel problem, described in [37], [38]. In the $K_t \times K_r$ compound X-channel, every transmitter has a message for every receiver. However, the channel of every receiver i can be one of some *finite* number J_i of states, and transmission has to account for all possible states. The optimal sum DoF in this problem is $\frac{K_t K_r}{K_t + K_r - 1}$, i.e., $\frac{1}{\frac{1}{K_t} + \frac{1}{K_r} - 1}$ per message, as stated in [37, Theorem 4].

If the receiver is able to decode its messages regardless of which of the J_i realizations the channel has taken, then this is equivalent to replacing the single receiver with J_i channel realizations by J_i different receivers with each a single possible channel realization, such that all J_i receivers want the same messages. This is exactly the problem statement we have at the physical layer. Our problem is therefore equivalent to a $K_t \times \tilde{N}$ compound X-channel with $|\mathcal{U}_n|$ channel realizations for every receiver n . Therefore, if $\tilde{R}_{\tilde{N}}$ denotes the rate of each message, and $\tilde{d}_{\tilde{N}}$ its DoF, then [37, Theorem 4] implies that the optimal DoF is

$$\tilde{d}_{\tilde{N}}(K_t, K_r) = \frac{1}{K_t + \tilde{N} - 1}. \quad (21)$$

B. Network Layer

Let the link load $\ell_{\tilde{N}}$ denote the size of each $V_{\mathcal{U}_n j}$ in units of files. The strategy at the network layer is straightforward. For every subset \mathcal{U}_n of users, each transmitter j sends the part of the file that they requested through $V_{\mathcal{U}_n j}$. Mathematically, we set

$$V_{\mathcal{U}_n j} = W_n^j,$$

for all $j = 1, \dots, K_t$ and n such that $\mathcal{U}_n \neq \emptyset$. This allows every user to decode its requested file. Since every multicast link $V_{\mathcal{U}_n j}$ carries one file part W_n^j , then the link load is

$$\ell_{\tilde{N}}(N, K_t, K_r, M_t, 0) = \frac{1}{K_t}. \quad (22)$$

C. Achievable End-to-End DoF

Note that the same $V_{\mathcal{U}_n j}$ has a size of $\tilde{R}_{\tilde{N}} T$ at the physical layer and $\ell_{\tilde{N}} F$ at the network layer. Since $F = RT$, we get $R = \tilde{R}_{\tilde{N}} / \ell_{\tilde{N}}$, and by combining that with (21) and (22), we achieve a DoF of

$$\frac{\tilde{d}_{\tilde{N}}(K_t, K_r)}{\ell_{\tilde{N}}(N, K_t, K_r, M_t, 0)} = \frac{K_t}{K_t + \tilde{N} - 1}.$$

In the worst case, the largest number of distinct files are requested, i.e., $\tilde{N} = \min\{K_r, N\}$. Therefore,

$$\text{DoF} \geq \frac{K_t}{K_t + \min\{K_r, N\} - 1},$$

when $M_r = 0$.

Since $1/\text{DoF}$ is convex in M_r , and $1/\text{DoF} = 0$ when $M_r = N$, then, for all intermediate values of M_r , we can achieve

$$\frac{1}{\text{DoF}} \leq \frac{K_t + \min\{K_r, N\} - 1}{K_t} \cdot \left(1 - \frac{M_r}{N}\right). \quad (23)$$

For any $M_r \in [0, N]$, we can choose whichever of the two separation interfaces yields the higher DoF. Therefore, combining (23) with (13) yields

$$\begin{aligned} \frac{1}{\text{DoF}} &\leq \min \left\{ \frac{K_t + \min\{K_r, N\} - 1}{K_t} \cdot \left(1 - \frac{\kappa}{K_r}\right), \frac{K_t - 1 + \frac{K_r}{\kappa+1}}{K_t} \cdot \left(1 - \frac{\kappa}{K_r}\right) \right\} \\ &= \frac{K_t + \min\left\{\frac{K_r}{\kappa+1}, N\right\} - 1}{K_t} \cdot \left(1 - \frac{\kappa}{K_r}\right), \end{aligned}$$

when $M_r = \kappa N/K_r$ with κ being an integer, and the lower convex envelope of these points for all $M_r \in [0, N]$. This concludes the achievability proof of Theorem 1.

APPENDIX B PROOF OF LEMMA 4

Let Γ be defined as in the statement of the lemma, and let $n \in \mathbb{N}$ be arbitrary. Define T_n as

$$T_n = \binom{K_r - 1}{\sigma - 1} [(n + 1)^\Gamma + (K_t - 1)n^\Gamma] + \binom{K_r - 1}{\sigma} (n + 1)^\Gamma.$$

We will show that a DoF of $\delta_j^{(n)}$ can be achieved for message V_{S_j} over a block length of T_n . We first describe how to (maximally) align the interference at each receiver, and then show that the receiver's desired messages are still decodable. The proof will rely on two lemmas from [4]: the alignment part will use [4, Lemma 2], while the decodability part will rely on [4, Lemma 1]. For ease of reference, we have rephrased the two lemmas in Appendix E as Lemmas 7 and 8, respectively.

Alignment: Describe each message V_{S_j} as a column vector of $(n + c_j)^\Gamma$ symbols $\mathbf{v}_{S_j} = [v_{S_j}^m]_{m=1}^{(n+c_j)^\Gamma}$, where c_j is as defined in the statement of the lemma. Each symbol is beamformed along a length- T_n vector $\mathbf{a}_{S_j}^m$, so that transmitter j sends the codeword

$$\mathbf{x}_j = \sum_{\mathcal{S}:|\mathcal{S}|=\sigma} \sum_{m=1}^{(n+c_j)^\Gamma} v_{S_j}^m \mathbf{a}_{S_j}^m,$$

over the block length T_n . We can alternatively combine all the $\mathbf{a}_{S_j}^m$ vectors into one matrix $\mathbf{A}_{S_j} = [\mathbf{a}_{S_j}^1, \dots, \mathbf{a}_{S_j}^{(n+c_j)^\Gamma}]$, and write

$$\mathbf{x}_j = \sum_{\mathcal{S}} \mathbf{A}_{S_j} \mathbf{v}_{S_j}.$$

Receiver i then observes

$$\mathbf{y}_i = \sum_{j=1}^{K_t} \mathbf{H}_{ij} \mathbf{x}_j + \mathbf{z}_i = \sum_{j=1}^{K_t} \mathbf{H}_{ij} \sum_{\mathcal{S}} \mathbf{A}_{S_j} \mathbf{v}_{S_j} + \mathbf{z}_i. \quad (24)$$

Recall that \mathbf{z}_i is the iid additive Gaussian unit-variance noise, and \mathbf{H}_{ij} is a $T_n \times T_n$ diagonal matrix representing the independent continuously-distributed channel coefficients over block length T_n , as defined

in Section II. In other words, the τ -th diagonal element of \mathbf{H}_{ij} is $h_{ij}(\tau)$. Moreover, the dimensions of \mathbf{A}_{S_j} are $T_n \times (n + c_j)^\Gamma$, and the length of \mathbf{v}_{S_j} is $(n + c_j)^\Gamma$.

In the expression for \mathbf{y}_i in (24), it will be convenient to separate the messages intended for i from the interfering messages,

$$\mathbf{y}_i = \sum_{S:i \in S} \sum_{j=1}^{K_t} \mathbf{H}_{ij} \mathbf{A}_{S_j} \mathbf{v}_{S_j} + \sum_{S:i \notin S} \left[\mathbf{H}_{i1} \mathbf{A}_{S_1} \mathbf{v}_{S_1} + \sum_{j=2}^{K_t} \mathbf{H}_{ij} \mathbf{A}_{S_j} \mathbf{v}_{S_j} \right] + \mathbf{z}_i.$$

Our goal is to collapse each term inside the second sum (i.e., for each S such that $i \notin S$) into a single subspace, namely the subspace spanned by $\mathbf{H}_{i1} \mathbf{A}_{S_1}$.⁵ This should be done for all $i \in \{1, \dots, K_r\}$. Specifically, we want to choose the \mathbf{A}_{S_j} 's such that they satisfy the following conditions almost surely:

$$\mathbf{H}_{ij} \mathbf{A}_{S_j} \prec \mathbf{H}_{i1} \mathbf{A}_{S_1}, \quad \forall i = 1, \dots, K_r, \quad \forall j = 2, \dots, K_t, \quad \forall S \text{ s.t. } i \notin S,$$

where $\mathbf{P} \prec \mathbf{Q}$ denotes that the vector space spanned by the columns of \mathbf{P} is a subspace of the one spanned by the columns of \mathbf{Q} .

First, we set $\mathbf{A}_{S_2} = \dots = \mathbf{A}_{S_{K_t}}$ for all subsets S . Thus we have reduced the problem to finding matrices \mathbf{A}_{S_1} and \mathbf{A}_{S_2} for all subsets S such that, almost surely,

$$\mathbf{H}_{i1}^{-1} \mathbf{H}_{ij} \mathbf{A}_{S_2} \prec \mathbf{A}_{S_1}, \quad \forall i \notin S, \quad \forall j = 2, \dots, K_t. \quad (25)$$

Note that \mathbf{H}_{i1}^{-1} exists almost surely since each diagonal element of \mathbf{H}_{i1} follows a continuous distribution and is thus non-zero with probability one.

For every S , the matrices \mathbf{A}_{S_1} and \mathbf{A}_{S_2} are constrained by a total of $(K_r - \sigma)(K_t - 1) = \Gamma$ subspace relations. We hence have Γ relations $\mathbf{G}_g \mathbf{A}_{S_2} \prec \mathbf{A}_{S_1}$, $g = 1, \dots, \Gamma$, where \mathbf{G}_g are $T_n \times T_n$ diagonal matrices. We can write all the diagonal elements of these matrices as forming the set

$$\mathcal{G} = \left\{ \frac{h_{ij}(\tau)}{h_{i1}(\tau)} : i \notin S, j \in \{2, \dots, K_t\}, \tau \in \{1, \dots, T_n\} \right\}.$$

Importantly, each element of \mathcal{G} follows a continuous distribution when conditioned on all the others, i.e.,

$$\frac{h_{ij}(\tau)}{h_{i1}(\tau)} \left| \left\{ \frac{h_{i'j'}(\tau')}{h_{i'1}(\tau')} : (i', j', \tau') \neq (i, j, \tau) \right\} \right. \sim \text{a continuous distribution.}$$

Furthermore, the dimensions of \mathbf{A}_{S_2} are $T_n \times n^\Gamma$, and the dimensions of \mathbf{A}_{S_1} are $T_n \times (n + 1)^\Gamma$, with $T_n > (n + 1)^\Gamma$.

Let P be some continuous probability distribution with a bounded support. For every $S \subseteq \{1, \dots, K_r\}$ such that $|\mathcal{S}| = \sigma$, we generate a $T_n \times 1$ column vector $\mathbf{b}_S = (b_S(1), \dots, b_S(T_n))^\top$, such that all the entries of all $\binom{K_r}{\sigma}$ vectors $\{\mathbf{b}_S\}_S$ are chosen iid from P . We can now invoke Lemma 7 in Appendix E to construct with probability one, for each S and using \mathbf{b}_S , full-rank matrices \mathbf{A}_{S_1} and \mathbf{A}_{S_2} that satisfy the subspace relations in (25) almost surely. Furthermore, the entries of the τ -th rows of both \mathbf{A}_{S_1} and \mathbf{A}_{S_2} are each a multi-variate monomial in the entries of the τ -th rows of \mathbf{b}_S and \mathbf{G}_g , $g = 1, \dots, \Gamma$, i.e.,

$$b_S(\tau) \quad \text{and} \quad \frac{h_{ij}(\tau)}{h_{i1}(\tau)}, i \notin S, j = 2, \dots, K_t. \quad (26)$$

Note that the monomial entries of \mathbf{A}_{S_1} are distinct; the same goes for the monomial entries of \mathbf{A}_{S_2} .

We have thus ensured alignment of the interfering messages. In the remainder of the proof we show that the desired messages are still almost surely decodable at every receiver.

⁵This is why we choose \mathbf{v}_{S_1} to be a longer vector than \mathbf{v}_{S_j} , $j \geq 2$: this choice makes $\mathbf{H}_{i1} \mathbf{A}_{S_1}$ the larger subspace, which allows us to align the other subspaces with it using Lemma 7.

Decodability: Recall that the total dimension of the vector space, i.e., the block length, is

$$T_n = \binom{K_r - 1}{\sigma - 1} [(n + 1)^\Gamma + (K_t - 1)n^\Gamma] + \binom{K_r - 1}{\sigma}.$$

Let us fix a receiver k . For this receiver, we have:

- $\binom{K_r - 1}{\sigma - 1}$ subspaces $\mathbf{H}_{k1}\mathbf{A}_{S1}$, $k \in \mathcal{S}$, of dimension $(n + 1)^\Gamma$ each, carrying the length- $(n + 1)^\Gamma$ vectors \mathbf{v}_{S1} that must be decoded by receiver k ;
- $(K_t - 1)\binom{K_r - 1}{\sigma - 1}$ subspaces $\mathbf{H}_{kj}\mathbf{A}_{S2}$, $k \in \mathcal{S}$ and $j = 2, \dots, K_t$, of dimension n^Γ each, carrying the length- n^Γ vectors \mathbf{v}_{S2} that must also be decoded by receiver k ;
- $\binom{K_r - 1}{\sigma}$ subspaces corresponding to $\mathbf{H}_{k1}\mathbf{A}_{S1}$, $k \notin \mathcal{S}$, of dimension $(n + 1)^\Gamma$ each, collectively carrying all the interference at receiver k .

Our goal is to show that the above subspaces are non-aligned for every receiver k , which implies that the desired messages are decodable with high probability for a large enough SNR.

Define matrices \mathbf{D}_k and \mathbf{I}_k representing the subspaces carrying the desired messages and the interference, respectively, by horizontally concatenating the subspaces:

$$\mathbf{D}_k = [\mathbf{H}_{k1}\mathbf{A}_{S1} \quad \mathbf{H}_{k2}\mathbf{A}_{S2} \quad \cdots \quad \mathbf{H}_{kK_t}\mathbf{A}_{S2}]_{S:k \in \mathcal{S}}; \quad (27a)$$

$$\mathbf{I}_k = [\mathbf{H}_{k1}\mathbf{A}_{S1}]_{S:k \notin \mathcal{S}}. \quad (27b)$$

Therefore, decodability at receiver k is ensured if the $T_n \times T_n$ matrix

$$\mathbf{\Psi}_k = [\mathbf{D}_k \quad \mathbf{I}_k]$$

is full rank almost surely. We prove that this is true with the help of Lemma 8 in Appendix E. To apply Lemma 8, we need to show that the following two conditions hold.

- 1) Two distinct rows of $\mathbf{\Psi}_k$ consist of monomials in disjoint sets of variables. In other words, the variables involved in the monomials of a specific row are exclusive to that row.
- 2) Within each row, each entry is a *unique* product of powers of the variables associated with that row.

To show that the first condition holds, consider the τ -th row of $\mathbf{\Psi}_k$. This row consists of monomial terms in the variables $\{b_S(\tau)\}_S$ and $\{h_{ij}(\tau)\}_{i,j}$. This is true because the τ -th row of any submatrix $\mathbf{H}_{kj}\mathbf{A}_{Sj}$ of $\mathbf{\Psi}_k$ is equal to $h_{kj}(\tau)$ multiplied by the τ -th row of \mathbf{A}_{Sj} , whose entries are monomials in the variables listed in (26). Therefore, the variables involved in a row of $\mathbf{\Psi}_k$ are exclusive to that row.

Before we prove that the second condition holds, we emphasize two remarks regarding the monomials that constitute the entries of the τ -th row of $\mathbf{\Psi}_k$.

Remark 1: All the entries in the τ -th row of submatrix $\mathbf{H}_{kj}\mathbf{A}_{Sj}$ are distinct monomials from one another. This is true because the τ -th row of $\mathbf{H}_{kj}\mathbf{A}_{Sj}$ is equal to the τ -th row of \mathbf{A}_{Sj} , whose entries are distinct monomials by construction (using Lemma 8), multiplied by $h_{kj}(\tau)$.

Remark 2: It follows from (26) that the entries of the τ -th row of submatrix $\mathbf{H}_{kj}\mathbf{A}_{Sj}$ are monomials in which only the variables in a set $\mathcal{B}_{Sj}(\tau)$ appear (with non-zero exponent), where $\mathcal{B}_{Sj}(\tau)$ obeys:

$$b_S(\tau) \in \mathcal{B}_{Sj}(\tau) \quad \text{and} \quad b_{S'}(\tau) \notin \mathcal{B}_{Sj}(\tau), \quad \forall S' \neq S; \quad (28a)$$

$$\text{and} \quad k \in \mathcal{S} \implies \mathcal{B}_{Sj}(\tau) = \{b_S(\tau), h_{kj}(\tau)\} \cup \{h_{i1}(\tau), \dots, h_{iK_t}(\tau) : i \notin \mathcal{S}\}. \quad (28b)$$

Note that when $k \notin \mathcal{S}$, we cannot be sure if $\mathcal{B}_{Sj}(\tau)$ contains $h_{kj}(\tau)$ because the latter is present in the monomials of both \mathbf{H}_{kj} and \mathbf{A}_{Sj} , and can hence be canceled out in their product.

Remark 1 states that monomials in the same submatrix are distinct. Therefore, all that remains is to show the same for monomials in the τ -th rows of different submatrices. However, by Remark 2 the same variables appear with non-zero exponent in *all* entries in the τ -th row of any submatrix (albeit with different powers). It is therefore sufficient to prove that the τ -th rows of two different submatrices $\mathbf{H}_{kj}\mathbf{A}_{Sj}$ and $\mathbf{H}_{kj'}\mathbf{A}_{Sj'}$ are functions of different sets of variables. Specifically, we show that there is a variable

TABLE I

Exponents of variables in monomials of the τ -th row of Ψ_k for an arbitrary receiver k . The subsets \mathcal{S}_1 , \mathcal{S}'_1 , \mathcal{S}_2 , and \mathcal{S}'_2 are arbitrary such that $k \in \mathcal{S}_1 \cap \mathcal{S}'_1$ and $k \notin \mathcal{S}_2 \cup \mathcal{S}'_2$. The transmitters j and j' are also arbitrary. A cell will contain a check mark (\checkmark) if the variable in the corresponding row appears with non-zero exponent in the monomials of the τ -th row of the submatrix in the corresponding column. The cell will be empty if the variable does not appear in those monomials. It will contain a question mark (?) if the variable may or may not appear. Not all variables and submatrices are shown; only a representative few are used. Finally, recall that $\mathbf{A}_{\mathcal{S}_j} = \mathbf{A}_{\mathcal{S}_2}$ for $j \geq 2$.

	\mathbf{I}_k		\mathbf{D}_k			
	$\mathbf{H}_{k1} \mathbf{A}_{\mathcal{S}_2 1}$	$\mathbf{H}_{k1} \mathbf{A}_{\mathcal{S}'_2 1}$	$\mathbf{H}_{kj} \mathbf{A}_{\mathcal{S}_1 j}$	$\mathbf{H}_{kj} \mathbf{A}_{\mathcal{S}'_1 j}$	$\mathbf{H}_{kj'} \mathbf{A}_{\mathcal{S}_1 j'}$	$\mathbf{H}_{kj'} \mathbf{A}_{\mathcal{S}'_1 j'}$
$b_{\mathcal{S}_1}(\tau)$			\checkmark		\checkmark	
$b_{\mathcal{S}'_1}(\tau)$				\checkmark		\checkmark
$b_{\mathcal{S}_2}(\tau)$	\checkmark					
$b_{\mathcal{S}'_2}(\tau)$		\checkmark				
$h_{kj}(\tau)$?	?	\checkmark	\checkmark		
$h_{kj'}(\tau)$?	?			\checkmark	\checkmark

that appears with non-zero exponent in all the entries of the τ -th row of one submatrix but in none of the entries of the τ -th row of the other. We will prove below that this claim is true using Remark 2, with the aid of Table I for visualization.

For convenience, define $\mathbf{r}_{\mathcal{S}_j}^\top$ to be the τ -th row of matrix $\mathbf{H}_{kj} \mathbf{A}_{\mathcal{S}_j}$, and similarly define $\mathbf{r}_{\mathcal{S}'_{j'}}^\top$ to be the τ -th row of matrix $\mathbf{H}_{kj'} \mathbf{A}_{\mathcal{S}'_{j'}}$. To show that the entries of these rows are monomial functions of distinct variables, we isolate two cases: case $\mathcal{S} \neq \mathcal{S}'$ and case $\mathcal{S} = \mathcal{S}'$, $j \neq j'$.

- 1) Suppose $\mathcal{S} \neq \mathcal{S}'$. Then, by (28a), $\mathbf{r}_{\mathcal{S}_j}^\top$ is a function of $b_{\mathcal{S}}(\tau)$ but not $b_{\mathcal{S}'}(\tau)$ while the opposite is true of $\mathbf{r}_{\mathcal{S}'_{j'}}^\top$.
- 2) Suppose now that $\mathcal{S} = \mathcal{S}'$ but $j \neq j'$. Crucially, two such matrices are relevant at receiver k only if $k \in \mathcal{S}$, as evidenced by (27). Therefore, by (28b), row $\mathbf{r}_{\mathcal{S}_j}^\top$ is a function of $h_{kj}(\tau)$ but not $h_{kj'}(\tau)$, and the reverse is true of $\mathbf{r}_{\mathcal{S}'_{j'}}^\top$.

Combining the above two points, it follows that the entries of the two rows are monomials in a different set of variables. We can conclude that the entries in the τ -th row of Ψ_k are distinct monomials, and specifically that the matrix Ψ_k is of the form seen in the statement of Lemma 8. Therefore, by Lemma 8, the matrix Ψ_k is full rank almost surely, and thus all receivers are able to decode their desired messages almost surely.

In conclusion: We were able to transmit all the messages $V_{\mathcal{S}_j}$, represented by length- $(n+c_j)^\Gamma$ vectors $\mathbf{v}_{\mathcal{S}_j}$, over a block length of T_n . Hence, the degrees of freedom achieved for each $V_{\mathcal{S}_j}$ is

$$\frac{(n+c_j)^\Gamma}{T_n} = \frac{(n+c_j)^\Gamma}{\binom{K_r-1}{\sigma-1} [(n+1)^\Gamma + (K_t-1)n^\Gamma] + \binom{K_r-1}{\sigma}(n+1)^\Gamma} = \delta_j^{(n)},$$

which concludes the proof of Lemma 4. ■

APPENDIX C

DETAILED CONVERSE PROOF OF THEOREM 1

A high-level overview of the converse proof of Theorem 1 was given in Section VI. In this appendix, we will give the rigorous proof. In particular, we will prove (16), i.e.,

$$\frac{1}{\text{DoF}} \geq \frac{d^{-1}(M_r)}{13.5}$$

by analyzing the four regimes described in (14).

Regime 0: $\min\{K_r, N\} \leq 12.5K_t$

In this regime, the number of transmitters is at least of the order of the total number of different requested files. As described in Section III-4, this implies that $1/\text{DoF} \approx 1 - \frac{M_r}{N}$. More precisely, by convexity of $d^{-1}(M_r)$ we have

$$d^{-1}(M_r) \leq d^{-1}(0) - \frac{d^{-1}(0) - d^{-1}(N)}{N - 0}(M_r - 0) = d^{-1}(0) \left(1 - \frac{M_r}{N}\right)$$

for all $M_r \in [0, N]$, where we have used that $d^{-1}(N) = 0$. Moreover, we have

$$d^{-1}(0) = \frac{K_t - 1 + \min\{K_r, N\}}{K_t} \leq \frac{K_t + 12.5K_t}{K_t} = 13.5,$$

which implies

$$d^{-1}(M_r) \leq 13.5 \left(1 - \frac{M_r}{N}\right). \quad (29)$$

We now invoke Lemma 5 with $s = 1$, yielding

$$\frac{1}{\text{DoF}} \geq 1 - \frac{M_r}{N} \stackrel{(a)}{\geq} \frac{d^{-1}(M_r)}{13.5} \quad \text{in Regime 0,} \quad (30)$$

where (a) follows from (29).

In all the following regimes, $\min\{K_r, N\} > 12.5K_t \geq 12.5$.

Regime 1: $0 \leq M_r \leq 1.1 \max\{1, N/K_r\}$

Since $d^{-1}(M_r)$ is non-increasing in M_r , we can upper-bound it by

$$d^{-1}(M_r) \leq d^{-1}(0) = \frac{K_t - 1 + \min\{K_r, N\}}{K_t} \leq \frac{\left(\frac{1}{12.5} + 1\right) \min\{K_r, N\}}{K_t} = \frac{13.5}{12.5} \cdot \frac{\min\{K_r, N\}}{K_t}. \quad (31)$$

Let us now use Lemma 5 with $s = \lfloor 0.275 \min\{K_r, N\} \rfloor \in \{1, \dots, \min\{K_r, N\}\}$. Then, using $\lfloor N/s \rfloor \geq N/s - 1$,

$$\begin{aligned} \frac{1}{\text{DoF}} &\geq \frac{s}{\min\{s, K_t\}} \left(1 - \frac{s}{1 - s/N} \cdot \frac{M_r}{N}\right) \\ &\geq \frac{1}{K_t} \left(s - \frac{s^2}{1 - s/N} \cdot \frac{1.1 \max\{1, N/K_r\}}{N}\right) \\ &= \frac{1}{K_t} \left(\lfloor 0.275 \min\{K_r, N\} \rfloor - \frac{[0.275 \min\{K_r, N\}]^2}{1 - [0.275 \min\{K_r, N\}]/N} \cdot \frac{1.1 \max\{1, N/K_r\}}{N}\right) \\ &\geq \frac{1}{K_t} \left(0.275 \min\{K_r, N\} - 1 - \frac{(0.275 \min\{K_r, N\})^2}{1 - 0.275} \cdot \frac{1.1 \max\{1, N/K_r\}}{N}\right) \\ &= \frac{1}{K_t} \left(0.275 \min\{K_r, N\} - 1 - \frac{(0.275 \min\{K_r, N\})^2}{0.725} \cdot \frac{1.1}{\min\{K_r, N\}}\right) \\ &= \frac{\min\{K_r, N\}}{K_t} \left(0.275 - \frac{1}{\min\{K_r, N\}} - \frac{(0.275)^2 \cdot 1.1}{0.725}\right) \\ &\geq \frac{\min\{K_r, N\}}{K_t} \left(0.275 - \frac{1}{12.5} - \frac{(0.275)^2 \cdot 1.1}{0.725}\right) \\ &\geq \frac{1}{12.5} \cdot \frac{\min\{K_r, N\}}{K_t} \\ &\stackrel{(a)}{\geq} \frac{d^{-1}(M_r)}{13.5} \quad \text{in Regime 1,} \end{aligned} \quad (32)$$

where (a) uses (31).

Regime 2: $1.1 \max\{1, N/K_r\} < M_r \leq 0.092N/K_t$

Let \tilde{M}_r be the largest integer multiple of N/K_r that is no greater than M_r , and define $\tilde{\kappa} = K_r \tilde{M}_r / N$. Note that $\tilde{\kappa}$ is an integer. Hence,

$$0 \leq M_r - N/K_r < \tilde{M}_r \leq M_r.$$

Since $d^{-1}(M_r)$ is non-increasing in M_r , we have:

$$\begin{aligned} d^{-1}(M_r) &\leq d^{-1}(\tilde{M}_r) \\ &= \frac{K_t - 1 + \min\{\frac{K_r}{\tilde{\kappa}+1}, N\}}{K_t} \cdot \left(1 - \frac{\tilde{\kappa}}{K_r}\right) \\ &\leq \frac{K_t + \frac{K_r}{\tilde{\kappa}+1}}{K_t} \\ &= 1 + \frac{K_r}{K_t(\tilde{\kappa} + 1)} \\ &= 1 + \frac{K_r}{K_t(K_r \tilde{M}_r / N + 1)} \end{aligned} \tag{33}$$

$$\begin{aligned} &\stackrel{(a)}{\leq} 1 + \frac{K_r}{K_t \cdot K_r M_r / N} \\ &= \frac{N}{K_t M_r} \left(\frac{K_t M_r}{N} + 1 \right) \\ &\stackrel{(b)}{\leq} 1.092 \frac{N}{K_t M_r}, \end{aligned} \tag{34}$$

where (a) uses $\tilde{M}_r > M_r - N/K_r$ and (b) follows from $M_r \leq 0.092N/K_t$.

We now invoke Lemma 5 with $s = \lfloor 0.3N/M_r \rfloor \in \{1, \dots, \min\{K_r, N\}\}$. Once again, we write:

$$\begin{aligned} \frac{1}{\text{DoF}} &\geq \frac{1}{\min\{s, K_t\}} \left(s - \frac{s^2}{1 - s/N} \cdot \frac{M_r}{N} \right) \\ &\geq \frac{1}{K_t} \left(0.3 \frac{N}{M_r} - 1 - \frac{(0.3N/M_r)^2}{1 - 0.3/M_r} \cdot \frac{M_r}{N} \right) \\ &\stackrel{(a)}{\geq} \frac{1}{K_t} \left(0.3 \frac{N}{M_r} - 1 - \frac{0.3^2 N/M_r}{1 - 0.3/1.1} \right) \\ &= \frac{N}{K_t M_r} \left(0.3 - \frac{M_r}{N} - \frac{0.3^2}{1 - 0.3/1.1} \right) \\ &\stackrel{(b)}{\geq} \frac{N}{K_t M_r} \left(0.3 - 0.092 - \frac{0.3^2}{1 - 0.3/1.1} \right) \\ &\geq \frac{1.092}{13.5} \cdot \frac{N}{K_t M_r} \\ &\stackrel{(c)}{\geq} \frac{d^{-1}(M_r)}{13.5} \quad \text{in Regime 2,} \end{aligned} \tag{35}$$

where (a) is due to $M_r \geq 1.1 \max\{1, N/K_r\} \geq 1.1$, (b) follows from $M < 0.092N/K_t \leq 0.092N$, and (c) uses (34).

Regime 3: $0.092N/K_t < M_r \leq N$

By the convexity of $d^{-1}(M_r)$, we have for all $M_r \in (0.092N/K_t, N]$,

$$\begin{aligned}
d^{-1}(M_r) &\leq d^{-1}(0.092N/K_t) - \frac{d^{-1}(0.092N/K_t) - d^{-1}(N)}{N - 0.092N/K_t} \cdot \left(M_r - 0.092\frac{N}{K_t}\right) \\
&\stackrel{(a)}{=} d^{-1}(0.092N/K_t) \left(1 - \frac{M_r - 0.092N/K_t}{N - 0.092N/K_t}\right) \\
&= d^{-1}(0.092N/K_t) \left(\frac{N - M_r}{N - 0.092N/K_t}\right) \\
&= \frac{d^{-1}(0.092N/K_t)}{1 - 0.092/K_t} \left(1 - \frac{M_r}{N}\right) \\
&\leq \frac{d^{-1}(0.092N/K_t)}{0.908} \left(1 - \frac{M_r}{N}\right), \tag{36}
\end{aligned}$$

where (a) uses that $d^{-1}(N) = 0$.

Let \tilde{M}_r be the largest integer multiple of N/K_r that is no greater than $0.092N/K_t$, and define $\tilde{\kappa} = K_r\tilde{M}_r/N$. Note that $\tilde{\kappa}$ is an integer. Then,

$$0 \stackrel{(a)}{\leq} 0.092N/K_t - N/K_r < \tilde{M}_r \leq 0.092N/K_t,$$

where (a) follows from (15). This implies that $\tilde{\kappa} + 1 \geq 0.092K_r/K_t$. Since $d^{-1}(M_r)$ is non-increasing,

$$\begin{aligned}
d^{-1}(0.092N/K_t) &\leq d^{-1}(\tilde{M}_r) \\
&\stackrel{(a)}{\leq} 1 + \frac{K_r}{K_t(\tilde{\kappa} + 1)} \\
&\leq 1 + \frac{K_r}{K_t \cdot 0.092K_r/K_t} \\
&= \frac{1.092}{0.092}, \tag{37}
\end{aligned}$$

where (a) follows the same steps that led to (33). Combining (37) with (36),

$$d^{-1}(M_r) \leq \frac{1.092}{0.092 \cdot 0.908} \left(1 - \frac{M_r}{N}\right) \leq 13.1 \left(1 - \frac{M_r}{N}\right). \tag{38}$$

By applying Lemma 5 with $s = 1$ again, we obtain

$$\frac{1}{\text{DoF}} \geq 1 - \frac{M_r}{N} \stackrel{(a)}{\geq} \frac{d^{-1}(M_r)}{13.5} \quad \text{in Regime 3,} \tag{39}$$

where (a) follows from (38).

Synthesis

The inequalities in (30), (32), (35), and (39) cover all possible regimes. Therefore, they together give $1/\text{DoF} \geq d^{-1}(M_r)/13.5$, or equivalently

$$\text{DoF} \leq 13.5 \cdot d(N, K_t, K_r, M_t, M_r),$$

for all $N, K_t, K_r, M_t \in [0, N]$, and $M_r \in [0, N]$. This concludes the converse proof of Theorem 1. \blacksquare

APPENDIX D
COMMUNICATION PROBLEM OUTER BOUNDS (CONVERSE PROOF OF THEOREM 2)

In Section IV we have described the separation architecture and the communication problem that emerges from it. We call the communication problem the multiple multicast X-channel problem. We state its DoF $\tilde{d}_\sigma^*(K_t, K_r)$ in Theorem 2 and show that it is achievable by interference alignment in Section V. In this appendix, we prove its optimality by deriving matching information-theoretic outer bounds. Specifically, we want to prove

$$\tilde{d}_\sigma^*(K_t, K_r) \leq \frac{1}{K_t \binom{K_r-1}{\sigma-1} + \binom{K_r-1}{\sigma}}, \quad (40)$$

for all K_t and K_r .

Consider the following subset of messages:

$$\mathcal{V} = \{V_{Sj} : 1 \in S \text{ or } j = 1\}. \quad (41)$$

It will be convenient to split \mathcal{V} into two disjoint parts,

$$\begin{aligned} \mathcal{V}^r &= \{V_{Sj} : 1 \in S, j \in \{1, \dots, K_t\}\}, \\ \mathcal{V}^t &= \{V_{Sj} : 1 \notin S, j = 1\}. \end{aligned}$$

In what follows, we will only focus on \mathcal{V} . All other messages, collectively denoted by

$$\bar{\mathcal{V}} = \{V_{Sj} : 1 \notin S \text{ and } j \neq 1\},$$

are made available to everyone by a genie. Furthermore, we lower the noise at receiver one by a fixed (non-vanishing) amount. Specifically, we replace \mathbf{y}_1 by

$$\tilde{\mathbf{y}}_1 = \sum_{j=1}^{K_t} \mathbf{H}_{1j} \mathbf{x}_j + \tilde{\mathbf{z}}_1,$$

where $\tilde{z}_1(\tau)$ are independent zero-mean Gaussian variables with variance

$$\text{var}(\tilde{z}_1(\tau)) = \min_{i=1, \dots, K_r} \left(\frac{h_{11}(\tau)}{h_{i1}(\tau)} \right)^2. \quad (42)$$

Note that $\text{var}(\tilde{z}_1(\tau)) \leq 1$ since we can set $i = 1$ in (42). Hence all the above changes can only improve capacity.

Consider all the receivers other than receiver one. Let a genie also give all of these receivers the subset \mathcal{V}^r . Again, this can only improve capacity. Hence, these receivers are given $\mathcal{V}^r \cup \bar{\mathcal{V}}$, which consists of all the messages that receiver one should decode, as well as all the messages of all transmitters other than transmitter one. Using this genie-given knowledge, every receiver can compute \mathbf{x}_j for all $j \neq 1$, and subtract all of them out of their output \mathbf{y}_i . In other words, receiver $i \neq 1$ can compute

$$\mathbf{y}'_i = \mathbf{y}_i - \sum_{j \neq 1} \mathbf{H}_{ij} \mathbf{x}_j = \mathbf{H}_{i1} \mathbf{x}_1 + \mathbf{z}_i. \quad (43)$$

Receiver i is still expected to decode some messages. Specifically, it must decode the subset of \mathcal{V}^t that is intended for it, i.e.,

$$\mathcal{V}_i^t = \{V_{Sj} : 1 \notin S \text{ and } i \in S \text{ and } j = 1\}.$$

Then, by Fano's inequality,

$$H(\mathcal{V}_i^t | \mathbf{y}_i, \mathcal{V}^r, \bar{\mathcal{V}}) \leq \varepsilon T. \quad (44)$$

We focus now on receiver one. From the problem requirements, it should be able to decode all of \mathcal{V}^r with high probability. After decoding \mathcal{V}^r , it has access to all the messages that receiver $i \neq 1$ has, and hence it too can subtract out \mathbf{x}_j , $j \neq 1$, from its output,

$$\tilde{\mathbf{y}}'_1 = \tilde{\mathbf{y}}_1 - \sum_{j \neq 1} \mathbf{H}_{1j} \mathbf{x}_j + \tilde{\mathbf{z}}_1 = \mathbf{H}_{11} \mathbf{x}_1 + \tilde{\mathbf{z}}_1.$$

Since \mathbf{H}_{11} is invertible almost surely, receiver one can then transform its output to get a statistical equivalent of the output of any other receiver. Indeed, it can compute

$$\check{\mathbf{y}}_1^{(i)} = \mathbf{H}_{i1} \mathbf{H}_{11}^{-1} \tilde{\mathbf{y}}'_1 = \mathbf{H}_{i1} \mathbf{x}_1 + (\mathbf{H}_{i1} \mathbf{H}_{11}^{-1}) \tilde{\mathbf{z}}_1 = \mathbf{H}_{i1} \mathbf{x}_1 + \check{\mathbf{z}}_1^{(i)}.$$

Since $\check{\mathbf{z}}_1^{(i)} = (\mathbf{H}_{i1} \mathbf{H}_{11}^{-1}) \tilde{\mathbf{z}}_1$ and the \mathbf{H}_{ij} matrices are diagonal, then the variables $\check{z}_1^{(i)}(\tau)$ are independent and have a variance of

$$\text{var} \left(\check{z}_1^{(i)}(\tau) \right) = \text{var} \left(\frac{h_{i1}(\tau)}{h_{11}(\tau)} \tilde{z}_1(\tau) \right) = \left(\frac{h_{i1}(\tau)}{h_{11}(\tau)} \right)^2 \cdot \text{var}(\tilde{z}_1(\tau)) \leq 1,$$

by (42). As a result, receiver one has at least as good a channel output as \mathbf{y}'_i in (43), and can thus decode anything that receiver i can. In particular, it can decode \mathcal{V}_i^t for all i , i.e.,

$$H(\mathcal{V}_i^t | \tilde{\mathbf{y}}_1, \mathcal{V}^r, \bar{\mathcal{V}}) \leq H(\mathcal{V}_i^t | \mathbf{y}_i, \mathcal{V}^r, \bar{\mathcal{V}}) \leq \varepsilon T, \quad (45)$$

using (44).

All of the above can be mathematically expressed in the following chain of inequalities, for any achievable \tilde{R}_σ .

$$\begin{aligned} |\mathcal{V}| \cdot \tilde{R}_\sigma T &= H(\mathcal{V}) \\ &\stackrel{(a)}{=} H(\mathcal{V} | \bar{\mathcal{V}}) \\ &= I(\mathcal{V}; \tilde{\mathbf{y}}_1 | \bar{\mathcal{V}}) + H(\mathcal{V} | \tilde{\mathbf{y}}_1, \bar{\mathcal{V}}) \\ &= I(\mathcal{V}; \tilde{\mathbf{y}}_1 | \bar{\mathcal{V}}) + H(\mathcal{V}^r | \tilde{\mathbf{y}}_1, \bar{\mathcal{V}}) + H(\mathcal{V}^t | \tilde{\mathbf{y}}_1, \mathcal{V}^r, \bar{\mathcal{V}}) \\ &\stackrel{(b)}{\leq} I(\mathcal{V}; \tilde{\mathbf{y}}_1 | \bar{\mathcal{V}}) + \varepsilon T + H(\mathcal{V}^t | \tilde{\mathbf{y}}_1, \mathcal{V}^r, \bar{\mathcal{V}}) \\ &\stackrel{(c)}{\leq} I(\mathcal{V}; \tilde{\mathbf{y}}_1 | \bar{\mathcal{V}}) + \varepsilon T + \sum_{i \neq 1} H(\mathcal{V}_i^t | \tilde{\mathbf{y}}_1, \mathcal{V}^r, \bar{\mathcal{V}}) \\ &\stackrel{(d)}{\leq} I(\mathcal{V}; \tilde{\mathbf{y}}_1 | \bar{\mathcal{V}}) + \varepsilon T + \sum_{i \neq 1} \varepsilon T \\ &\stackrel{(e)}{\leq} I(\mathbf{x}_1, \dots, \mathbf{x}_{K_t}; \tilde{\mathbf{y}}_1 | \bar{\mathcal{V}}) + K_r \varepsilon T \\ &\stackrel{(f)}{\leq} T \cdot \left(\frac{1}{2} \log \text{SNR} + o(\log \text{SNR}) \right) + K_r \varepsilon T. \end{aligned}$$

In the above,

- (a) is due to the independence of the messages;
- (b) uses Fano's inequality for receiver one;
- (c) follows from observing that $\mathcal{V}^t = \bigcup_{i \neq 1} \mathcal{V}_i^t$;
- (d) uses (45);
- (e) is due to the data processing inequality; and
- (f) is the MAC channel bound.

By taking $T \rightarrow \infty$ and $\varepsilon \rightarrow 0$, as well as $\text{SNR} \rightarrow \infty$, we obtain

$$\tilde{d}_\sigma^*(K_t, K_r) \leq \frac{1}{|\mathcal{V}|} = \frac{1}{|\mathcal{V}^r| + |\mathcal{V}^t|} = \frac{1}{K_t \binom{K_r-1}{\sigma-1} + \binom{K_r-1}{\sigma}},$$

thus proving (40) and the converse part of Theorem 2. ■

APPENDIX E
LEMMAS FROM [4]

In our interference alignment strategy, we use two crucial lemmas from [4]. We present them here for ease of reference.

Lemma 7 (from [4, Lemma 2]). *Let $\mathbf{G}_1, \dots, \mathbf{G}_\Gamma$ be $T \times T$ diagonal matrices, such that $G_g(\tau)$, the τ -th diagonal entry of \mathbf{G}_g , follows a continuous distribution when conditioned on all other entries of all matrices. Also let \mathbf{b} be a column vector whose entries $b(\tau)$ are drawn iid from some continuous distribution, independently of $\mathbf{G}_1, \dots, \mathbf{G}_\Gamma$. Then, almost surely for any integer n such that $T > (n+1)^\Gamma$, there exist matrices \mathbf{A}_1 and \mathbf{A}_2 , of sizes $T \times (n+1)^\Gamma$ and $T \times n^\Gamma$ respectively, such that:*

- *Every entry in the τ -th row of \mathbf{A}_1 is a unique multi-variate monomial function of $b(\tau)$ and $G_g(\tau)$ for all g ($b(\tau)$ and $G_g(\tau)$ appear with non-zero exponents in every entry), and the same is true for \mathbf{A}_2 ;⁶ and*
- *The matrices satisfy the following conditions almost surely,*

$$\mathbf{G}_g \mathbf{A}_2 \prec \mathbf{A}_1, \quad \forall g = 1, \dots, \Gamma,$$

where $\mathbf{P} \prec \mathbf{Q}$ means that the span of the columns of \mathbf{P} is a subspace of the space spanned by the columns of \mathbf{Q} .

Lemma 8 (from [4, Lemma 1]). *Let $x_i^{(k)}$, $i = 1, \dots, T$ and $k = 1, \dots, K$, be random variables such that each follows a continuous distribution when conditioned on all other variables. Let Ψ be a $T \times T$ square matrix with entries ψ_{ij} such that*

$$\psi_{ij} = \prod_{k=1}^K \left(x_i^{(k)} \right)^{p_{ij}^{(k)}},$$

where $p_{ij}^{(k)}$ are integers such that

$$\left(p_{ij}^{(1)}, \dots, p_{ij}^{(K)} \right) \neq \left(p_{ij'}^{(1)}, \dots, p_{ij'}^{(K)} \right),$$

for all i, j, j' such that $j \neq j'$. In other words, the entries ψ_{ij} are distinct monomials in the variables $x_i^{(k)}$. Then, the matrix Ψ is almost surely full rank.

APPENDIX F
PROOF OF THEOREM 6

Theorem 6 gives an improved achievable DoF for the 2×2 cache-aided interference network. In this appendix, we prove this result by describing and analyzing the interference-extraction scheme introduced in Section VII and illustrated in Fig. 10, which achieves this DoF.

We describe the scheme in two steps. First, we focus on the physical layer to show how more information can be extracted from the aligned interference at the receivers. Second, we show how this additional information can be used at the network layer to achieve the inverse DoF in Theorem 6.

A. Physical Layer

In order to describe the interference-extraction scheme, let us first revisit the original separation architecture used when $M_r = 0$. The message set used for this case is the one where every transmitter has a message for every individual receiver, i.e., the unicast X-channel message set. In order to achieve the optimal communication DoF of $1/3$ per message, at every receiver, the two messages intended for the other receiver are aligned in the same subspace. Let us study this alignment more carefully.

⁶To clarify: a monomial could appear in both matrices \mathbf{A}_1 and \mathbf{A}_2 , but never twice in the same matrix.

Let V_{ij} be the message intended for receiver i from transmitter j . Represent every message V_{ij} by a scalar v_{ij} , called a stream. By taking a block length of 3 and by beamforming message V_{ij} along some direction \mathbf{a}_{ij} , we get channel inputs

$$\mathbf{x}_j = \mathbf{a}_{1j}v_{1j} + \mathbf{a}_{2j}v_{2j},$$

and channel outputs

$$\mathbf{y}_1 = \mathbf{H}_{11}(\mathbf{a}_{11}v_{11} + \mathbf{a}_{21}v_{21}) + \mathbf{H}_{12}(\mathbf{a}_{12}v_{12} + \mathbf{a}_{22}v_{22}) + \mathbf{z}_1; \quad (46a)$$

$$\mathbf{y}_2 = \mathbf{H}_{21}(\mathbf{a}_{11}v_{11} + \mathbf{a}_{21}v_{21}) + \mathbf{H}_{22}(\mathbf{a}_{12}v_{12} + \mathbf{a}_{22}v_{22}) + \mathbf{z}_2. \quad (46b)$$

Lemma 9. *We can choose the \mathbf{a}_{ij} 's in (46) such that*

$$\mathbf{y}_1 = \mathbf{\Psi}_1 \begin{bmatrix} v_{11} \\ v_{12} \\ v_{21} + v_{22} \end{bmatrix} + \mathbf{z}_1;$$

$$\mathbf{y}_2 = \mathbf{\Psi}_2 \begin{bmatrix} v_{21} \\ v_{22} \\ v_{11} + v_{12} \end{bmatrix} + \mathbf{z}_2,$$

where the 3×3 matrices $\mathbf{\Psi}_i$ are full-rank almost surely.

Proof: Recall that the \mathbf{H}_{ij} 's are 3×3 diagonal matrices whose τ -th diagonal element is $h_{ij}(\tau)$. Also recall that these $h_{ij}(\tau)$ are independent and continuously distributed, which implies that \mathbf{H}_{ij} is invertible almost surely. Assume this invertibility is the case in the following.

Choose the \mathbf{a}_{ij} vectors as:

$$\mathbf{a}_{11} = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}; \quad \mathbf{a}_{12} = \mathbf{H}_{22}^{-1}\mathbf{H}_{21}\mathbf{a}_{11};$$

$$\mathbf{a}_{21} = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}; \quad \mathbf{a}_{22} = \mathbf{H}_{12}^{-1}\mathbf{H}_{11}\mathbf{a}_{21}.$$

From (46), the received signals are then

$$\begin{aligned} \mathbf{y}_1 &= \mathbf{H}_{11}\mathbf{a}_{11}v_{11} + \mathbf{H}_{12}\mathbf{H}_{22}^{-1}\mathbf{H}_{21}\mathbf{a}_{11}v_{12} + \mathbf{H}_{11}\mathbf{a}_{21}(v_{21} + v_{22}) + \mathbf{z}_1 \\ &= [\mathbf{H}_{11}\mathbf{a}_{11} \quad \mathbf{H}_{12}\mathbf{H}_{22}^{-1}\mathbf{H}_{21}\mathbf{a}_{11} \quad \mathbf{H}_{11}\mathbf{a}_{21}] \begin{bmatrix} v_{11} \\ v_{12} \\ v_{21} + v_{22} \end{bmatrix} + \mathbf{z}_1 \\ &= \underbrace{\begin{bmatrix} h_{11}(1) & \frac{h_{12}(1)h_{21}(1)}{h_{22}(1)} & h_{11}(1) \\ h_{11}(2) & \frac{h_{12}(2)h_{21}(2)}{h_{22}(2)} & 0 \\ 0 & 0 & h_{11}(3) \end{bmatrix}}_{\mathbf{\Psi}_1} \begin{bmatrix} v_{11} \\ v_{12} \\ v_{21} + v_{22} \end{bmatrix} + \mathbf{z}_1, \end{aligned}$$

and, in a similar way,

$$\mathbf{y}_2 = \underbrace{\begin{bmatrix} h_{21}(1) & \frac{h_{22}(1)h_{11}(1)}{h_{12}(1)} & h_{21}(1) \\ 0 & 0 & h_{21}(2) \\ h_{21}(3) & \frac{h_{22}(3)h_{11}(3)}{h_{12}(3)} & 0 \end{bmatrix}}_{\mathbf{\Psi}_2} \begin{bmatrix} v_{21} \\ v_{22} \\ v_{11} + v_{12} \end{bmatrix} + \mathbf{z}_2.$$

Since the $h_{ij}(\tau)$ are independent continuously distributed variables, then the matrices $\mathbf{\Psi}_1$ and $\mathbf{\Psi}_2$ are full-rank almost surely. \blacksquare

Notice from Lemma 9 that each receiver can recover, in addition to its intended streams, the sum of the streams intended for the other receiver. By using a linear outer code over some finite field, we can ensure that obtaining the sum of two streams, e.g., $v_{21} + v_{22}$, yields the sum of the two corresponding messages, e.g., $V_{21} \oplus V_{22}$, where \oplus indicates addition over the finite field. For simplicity, we assume that this field is $\text{GF}(2)$, although any finite field gives the same result. In other words, receiver one can decode V_{11} , V_{12} , and $(V_{21} \oplus V_{22})$, and receiver two can decode V_{21} , V_{22} , and $(V_{11} \oplus V_{12})$. Therefore, for the same per-message DoF of $1/3$, we get the linear combinations of the unintended messages for free.

B. Network Layer

Fig. 10 illustrates the interface between the physical and network layers resulting from the decoding of the aligned interference at each receiver. This aligned interference, while available for free (no drawbacks in the communication DoF at the physical layer), becomes useful when the receiver memory is non-zero. It provides a middle ground between pure unicast messages (as is done at $M_r = 0$) and pure broadcast messages (which we use when $M_r = 1$).

Let ℓ denote the link load, i.e., the size of each message V_{ij} , and let $L = 4\ell$ be the sum network load. For this specific separation architecture, we denote by $L^*(M_r)$ the smallest sum network load as a function of receiver memory M_r , and by $\ell^* = L^*/4$ the smallest individual link load. Since each message V_{ij} (link) can be communicated across the physical layer using a DoF of $1/3$ by Lemma 9, then we can achieve an end-to-end DoF of

$$\frac{1}{\text{DoF}} \leq \frac{\ell^*}{1/3} = \frac{3}{4}L^*. \quad (47)$$

Theorem 6 follows directly from combining (47) with the following lemma.

Lemma 10. *For the separation architecture illustrated in Fig. 10, we can achieve the following sum network load:*

$$L^*(M_r) \leq \max \left\{ 2 - 2M_r, \frac{12}{7} - \frac{8}{7}M_r, \frac{4}{3} - \frac{2}{3}M_r \right\},$$

for $M_r \in [0, 2]$.


Proof: In order to prove Lemma 10, it suffices to look at the following four (M_r, L) corner points, as the rest can be achieved using time- and memory-sharing:

$$(0, 2), \quad (1/3, 4/3), \quad (4/5, 4/5), \quad (2, 0).$$

The fourth corner point is trivial since $M_r = 2$ implies each user can cache the entire library, and hence there is no need to transmit any information across the network. The first corner point can be achieved by ignoring the aligned interference messages, which reduces to the original strategy. Therefore, we only need to show the achievability of the second and third corner points. For convenience, we will call the two files in the content library A and B .


Achieving point $(M_r, L) = (1/3, 4/3)$: When $M_r = 1/3$, we split each file into three equal parts, labeled $A = (A_1, A_2, A_3)$ and $B = (B_1, B_2, B_3)$.

Table II shows the placement and delivery phases, for all possible user requests, and Fig. 11 illustrates the strategy when the demands are (A, B) . Notice that the transmitter caches hold exactly one file each (thus $M_t = 1$), the receivers cache one third of a file each ($M_r = 1/3$). Furthermore, the messages V_{ij} each carry the equivalent of one third of a file, which implies that $\ell = 1/3$ is achieved, or, equivalently, a sum network load of $L = 4\ell = 4/3$.

Fig. 11. Strategy for $M_r = 1/3$, when the user requests are (A, B) .TABLE II
Achievable strategy for $M_r = 1/3$.

Cache	Content				Rx
Tx 1	$A_3, B_1 \oplus B_3, B_2 \oplus B_3$				N/A
Tx 2	$B_3, A_1 \oplus A_3, A_2 \oplus A_3$				N/A
Rx 1	$A_1 \oplus B_1$				1
Rx 2	$A_2 \oplus B_2$				2

Message	Demands (Rx1, Rx2)				Rx
	(A, A)	(A, B)	(B, A)	(B, B)	
V_{11}	A_3	A_3	$B_2 \oplus B_3$	$B_1 \oplus B_3$	1
V_{21}	A_3	$B_1 \oplus B_3$	A_3	$B_2 \oplus B_3$	2
V_{12}	$A_1 \oplus A_3$	$A_2 \oplus A_3$	B_3	B_3	1
V_{22}	$A_2 \oplus A_3$	B_3	$A_1 \oplus A_3$	B_3	2
$V_{21} \oplus V_{22}$	A_2	B_1	A_1	B_2	1
$V_{11} \oplus V_{12}$	A_1	A_2	B_2	B_1	2

Fig. 12. Strategy for $M_r = 4/5$, when the user requests are (A, B) .TABLE III
Achievable strategy for $M_r = 4/5$.

Cache	Content				Rx
Tx 1	$A_5, B_5 \oplus S_1, B_5 \oplus S_2, B_5 \oplus S_3, B_5 \oplus S_4$				N/A
Tx 2	$B_5, A_5 \oplus T_1, A_5 \oplus T_2, A_5 \oplus T_3, A_5 \oplus T_4$				N/A
Rx 1	A_1, A_2, B_1, B_2				1
Rx 2	A_3, A_4, B_3, B_4				2

Message	Demands (Rx1, Rx2)				Rx
	(A, A)	(A, B)	(B, A)	(B, B)	
V_{11}	A_5	A_5	$B_5 \oplus S_2$	$B_5 \oplus S_3$	1
V_{21}	A_5	$B_5 \oplus S_1$	A_5	$B_5 \oplus S_4$	2
V_{12}	$A_5 \oplus T_1$	$A_5 \oplus T_3$	B_5	B_5	1
V_{22}	$A_5 \oplus T_2$	B_5	$A_5 \oplus T_4$	B_5	2
$V_{21} \oplus V_{22}$	T_2	S_1	T_4	S_4	1
$V_{11} \oplus V_{12}$	T_1	T_3	S_2	S_3	2

Achieving point $(M_r, L) = (4/5, 4/5)$: When $M_r = 4/5$, we split each file into five equal parts, labeled $A = (A_1, \dots, A_5)$ and $B = (B_1, \dots, B_5)$. For convenience, we define

$$\begin{aligned} S_1 &= B_2 \oplus A_4, & S_2 &= A_1 \oplus B_3, & S_3 &= B_1 \oplus B_3, & S_4 &= B_2 \oplus B_4, \\ T_1 &= A_1 \oplus A_3, & T_2 &= A_2 \oplus A_4, & T_3 &= B_1 \oplus A_3, & T_4 &= A_2 \oplus B_4, \end{aligned}$$

and write $\mathcal{S} = \{S_1, S_2, S_3, S_4\}$ and $\mathcal{T} = \{T_1, T_2, T_3, T_4\}$.

Table III shows the placement and delivery phases, for all possible user requests, and Fig. 12 illustrates the strategy when the demands are (A, B) . Notice that the transmitter caches hold exactly one file each (thus $M_t = 1$), the receivers cache four fifths of a file each ($M_r = 4/5$). Furthermore, the messages V_{ij} each carry the equivalent of one fifth of a file, which implies that $\ell = 1/5$ is achieved, or, equivalently, a sum network load of $L = 4\ell = 4/5$.

By achieving all four corner points, we have proved Lemma 10. ■

C. Optimality Within the Considered Separation Architecture

Within the separation architecture considered throughout this appendix and Section VII, i.e., the one illustrated in Fig. 10, we can show that the network-layer scheme is in fact exactly optimal. Specifically, the sum network load achieved in Lemma 10 is optimal. This is summarized in the following result.

Proposition 11. *For all M_r , the optimal sum network load must satisfy*

$$L^*(M_r) \geq \max \left\{ 2 - 2M_r, \frac{12}{7} - \frac{8}{7}M_r, \frac{4}{3} - \frac{2}{3}M_r \right\}.$$

While this does not contribute to the main result in Theorem 6, it does reinforce it by showing that this is the best we can do within this separation architecture.

Proof: For the proof, it is more convenient to write the outer bounds in terms of the individual link load $\ell^* = L^*/4$. Therefore, we will prove Proposition 11 by proving the following three inequalities (which together constitute an equivalent result):

$$\begin{aligned} 4\ell^* + 2M_r &\geq 2; \\ 7\ell^* + 2M_r &\geq 3; \\ 6\ell^* + M_r &\geq 2. \end{aligned}$$

In the following, we refer to the two files as A and B . Let the cache contents of receivers one and two be Q_1 and Q_2 , respectively. We also write V_{ij}^{ST} to denote the message V_{ij} when user one has requested file S and user two has requested file T , where $S, T \in \{A, B\}$. Furthermore, we use \mathcal{V}^{ST} to refer to all four messages when the requests are S and T , and write \mathcal{Y}_i^{ST} to denote the three outputs at receiver $i \in \{1, 2\}$ when the requests are S and T . Therefore,

$$\begin{aligned} \mathcal{V}^{ST} &= (V_{11}^{ST}, V_{12}^{ST}, V_{21}^{ST}, V_{22}^{ST}); \\ \mathcal{Y}_1^{ST} &= (V_{11}^{ST}, V_{12}^{ST}, V_{21}^{ST} \oplus V_{22}^{ST}); \\ \mathcal{Y}_2^{ST} &= (V_{21}^{ST}, V_{22}^{ST}, V_{11}^{ST} \oplus V_{12}^{ST}). \end{aligned}$$

We will next prove each of the three inequalities.

First inequality:

$$\begin{aligned}
(4\ell^* + 2M_r)F &\geq H(Q_1, Q_2, \mathcal{V}^{AB}) \\
&= H(Q_1, Q_2, \mathcal{V}^{AB} | A, B) + I(A, B; Q_1, Q_2, \mathcal{V}^{AB}) \\
&= H(Q_1, Q_2, \mathcal{V}^{AB} | A, B) + H(A, B) - H(A, B | Q_1, Q_2, \mathcal{V}^{AB}) \\
&\stackrel{(a)}{\geq} H(A, B) - \varepsilon F \\
&= 2F - \varepsilon F,
\end{aligned}$$

where (a) is due to Fano's inequality. Therefore,

$$4\ell^* + 2M_r \geq 2.$$

Second inequality:

$$\begin{aligned}
(7\ell^* + 2M_r)F &\geq H(Q_1, \mathcal{Y}_1^{AB}) + H(Q_2, \mathcal{V}^{BA}) \\
&\stackrel{(a)}{\geq} H(Q_1, \mathcal{Y}_1^{AB} | A) + H(Q_2, \mathcal{V}^{BA} | A) + 2H(A) - 2\varepsilon F \\
&\geq H(Q_1, Q_2, \mathcal{Y}_1^{AB}, \mathcal{V}^{BA} | A) + 2H(A) - 2\varepsilon F \\
&\stackrel{(b)}{\geq} H(Q_1, Q_2, \mathcal{Y}_1^{AB}, \mathcal{V}^{BA} | A, B) + H(B) + 2H(A) - 3\varepsilon F \\
&\geq 3F - 3\varepsilon F,
\end{aligned}$$

where (a) and (b) once again follow from Fano's inequality. Therefore,

$$7\ell^* + 2M_r \geq 3.$$

Third inequality:

$$\begin{aligned}
(6\ell^* + M_r)F &\geq H(Q_1, \mathcal{Y}_1^{AA}, \mathcal{Y}_2^{BB}) \\
&\stackrel{(a)}{\geq} H(Q_1, \mathcal{Y}_1^{AA}, \mathcal{Y}_2^{BB} | A, B) + H(A, B) - \varepsilon F \\
&\geq 2F - \varepsilon F,
\end{aligned}$$

where (a) is again due to Fano's inequality. Therefore,

$$6\ell^* + M_r \geq 2.$$

This concludes the proof of Proposition 11. ■

REFERENCES

- [1] V. Jacobson, D. K. Smetters, J. D. Thornton, M. Plass, N. Briggs, and R. Braynard, "Networking named content," *Comm. ACM*, vol. 55, no. 1, pp. 117–124, Jan. 2012.
- [2] M. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 2856–2867, May 2014.
- [3] —, "Cache-aided interference channels," in *Proc. IEEE ISIT*, June 2015, pp. 809–813.
- [4] V. R. Cadambe and S. A. Jafar, "Interference alignment and the degrees of freedom of wireless X networks," *IEEE Trans. Inf. Theory*, vol. 55, no. 9, pp. 3893–3908, Sept 2009.
- [5] D. Wessels, *Web Caching*. O'Reilly Media, Inc., 2001.
- [6] M. R. Korupolu, C. G. Plaxton, and R. Rajaraman, "Placement algorithms for hierarchical cooperative caching," in *Proc. ACM-SIAM SODA*, 1999, pp. 586–595.
- [7] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *Proc. IEEE INFOCOM*, 2010, pp. 1478–1486.
- [8] B. Tan and L. Massoulié, "Optimal content placement for peer-to-peer video-on-demand systems," *IEEE/ACM Trans. Netw.*, vol. 21, no. 2, pp. 566–579, Apr. 2013.
- [9] J. Llorca, A. M. Tulino, K. Guan, J. Esteban, M. Varvello, N. Choi, and D. C. Kilper, "Dynamic in-network caching for energy efficient content delivery," in *Proc. IEEE INFOCOM*, 2013, pp. 245–249.
- [10] A. Wolman, G. M. Voelker, N. Sharma, N. Cardwell, A. Karlin, and H. M. Levy, "On the scale and performance of cooperative web proxy caching," in *Proc. ACM SOSP*, 1999, pp. 16–31.

- [11] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and Zipf-like distributions: evidence and implications," in *Proc. IEEE INFOCOM*, 1999, pp. 126–134.
- [12] D. Applegate, A. Archer, V. Gopalakrishnan, S. Lee, and K. K. Ramakrishnan, "Optimal content placement for a large-scale VoD system," in *Proc. ACM CoNEXT*, 2010, pp. 4:1–4:12.
- [13] R. Pedarsani, M. A. Maddah-Ali, and U. Niesen, "Online coded caching," *IEEE/ACM Trans. Netw.*, vol. 24, no. 2, pp. 836–845, April 2016.
- [14] U. Niesen and M. A. Maddah-Ali, "Coded caching for delay-sensitive content," in *Proc. IEEE ICC*, June 2015, pp. 5559–5564.
- [15] S. Wang, W. Li, X. Tian, and H. Liu, "Coded caching with heterogeneous cache sizes," *arXiv:1504.01123v3 [cs.IT]*, Aug. 2015.
- [16] J. Zhang, X. Lin, C.-C. Wang, and X. Wang, "Coded caching for files with distinct file sizes," in *Proc. IEEE ISIT*, June 2015, pp. 1686–1690.
- [17] H. Ghasemi and A. Ramamoorthy, "Improved lower bounds for coded caching," in *Proc. IEEE ISIT*, June 2015, pp. 1696–1700.
- [18] A. Sengupta, R. Tandon, and T. Clancy, "Improved approximation of storage-rate tradeoff for caching via new outer bounds," in *Proc. IEEE ISIT*, June 2015, pp. 1691–1695.
- [19] M. Ji, G. Caire, and A. Molisch, "Wireless device-to-device caching networks: Basic principles and system performance," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 1, pp. 176–189, Jan 2016.
- [20] S. P. Shariatpanahi, A. S. Motahari, and B. H. Khalaj, "Multi-server coded caching," *arXiv:1503.00265v1 [cs.IT]*, Mar. 2015.
- [21] N. Golrezaei, K. Shanmugam, A. Dimakis, A. Molisch, and G. Caire, "Femtocaching: Wireless video content delivery through distributed caching helpers," in *Proc. IEEE INFOCOM*, March 2012, pp. 1107–1115.
- [22] J. Hachem, N. Karamchandani, and S. Diggavi, "Multi-level coded caching," in *Proc. IEEE ISIT*, Jun. 2014.
- [23] —, "Content caching and delivery over heterogeneous wireless networks," in *Proc. IEEE INFOCOM*, April 2015, pp. 756–764.
- [24] N. Karamchandani, U. Niesen, M. A. Maddah-Ali, and S. N. Diggavi, "Hierarchical coded caching," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3212–3229, June 2016.
- [25] U. Niesen and M. A. Maddah-Ali, "Coded caching with nonuniform demands," in *Proc. IEEE INFOCOM WKSHPs*, Apr. 2014, pp. 221–226.
- [26] M. Ji, A. M. Tulino, J. Llorca, and G. Caire, "On the average performance of caching and coded multicasting with random demands," in *Proc. IEEE ISWCS*, Aug. 2014.
- [27] J. Zhang, X. Lin, and X. Wang, "Coded caching under arbitrary popularity distributions," in *Proc. ITA*, Feb. 2015.
- [28] M. Ji, A. M. Tulino, J. Llorca, and G. Caire, "Caching-aided coded multicasting with multiple random requests," *arXiv:1511.07542 [cs.IT]*, Nov. 2015.
- [29] S. Gitzenis, G. S. Paschos, and L. Tassiulas, "Asymptotic laws for joint content replication and delivery in wireless networks," *IEEE Trans. Inf. Theory*, vol. 59, no. 5, pp. 2760–2776, May 2013.
- [30] S. Ioannidis, L. Massoulié, and A. Chaintreau, "Distributed caching over heterogeneous mobile networks," in *Proc. ACM SIGMETRICS*, 2010, pp. 311–322.
- [31] A. Sengupta, R. Tandon, and O. Simeone, "Cloud and cache-aided wireless networks: Fundamental latency trade-offs," in *Proc. IEEE ISIT*, Jul. 2016.
- [32] E. Altman, K. Avrachenkov, and J. Goseling, "Coding for caches in the plane," *arXiv:1309.0604 [cs.NI]*, Sep. 2013.
- [33] J. Y. Yang and B. Hajek, "Single video performance analysis for video-on-demand systems," *arXiv:1307.0849 [cs.NI]*, Jul. 2013.
- [34] N. Naderializadeh, M. A. Maddah-Ali, and A. S. Avestimehr, "Fundamental limits of cache-aided interference management," in *Proc. IEEE ISIT*, Jul. 2016.
- [35] F. Xu, M. Tao, and K. Liu, "Fundamental tradeoff between storage and latency in cache-aided wireless interference networks," in *Proc. IEEE ISIT*, Jul. 2016.
- [36] J. Hachem, U. Niesen, and S. Diggavi, "A layered caching architecture for the interference channel," in *Proc. IEEE ISIT*, Jul. 2016.
- [37] M. A. Maddah-Ali, "The degrees of freedom of the compound MIMO broadcast channels with finite states," *arXiv:0909.5006 [cs.IT]*, Sep. 2009.
- [38] T. Gou, S. A. Jafar, and C. Wang, "On the degrees of freedom of finite state compound wireless networks," *IEEE Trans. Inf. Theory*, vol. 57, no. 6, pp. 3286–3308, June 2011.