Ergodic Interference Alignment

Bobak Nazer and Michael Gastpar
University of California, Berkeley
EECS Department
Berkeley, CA 94720, USA
Email: {bobak, gastpar}@eecs.berkeley.edu

Syed Ali Jafar
University of California, Irvine
EECS Department
Irvine, CA, USA
Email: syed@eecs.uci.edu

Sriram Vishwanath
University of Texas, Austin
ECE Department
Austin, TX, USA
Email: sriram@ece.utexas.edu

***Abstract*—Consider a K-user interference channel with time-varying fading. At any particular time, each receiver will see a signal from most transmitters. The standard approach to such a scenario results in each transmitter-receiver pair achieving a rate proportional to $\frac{1}{K}$ the single user rate. However, given two well chosen time indices, the channel coefficients from interfering users can be made to exactly cancel. By adding up these two signals, the receiver can see an interference-free version of the desired transmission. We show that this technique allows each user to achieve at least half its interference-free ergodic capacity at any SNR. Prior work was only able to show that half the interference-free rate was achievable as the SNR tended to infinity. We examine a finite field channel model and a Gaussian channel model. In both cases, the achievable rate region has a simple description and, in the finite field case, we prove it is the ergodic capacity region.**

I. INTRODUCTION

The interference channel is one of the fundamental building blocks of wireless networks. Following several recent advances, the capacity region of the classical two-user Gaussian interference channel is known exactly for some interesting special cases (e.g. very weak or strong interference), and approximately (within one bit) for all channel conditions [1]. There is also increasing interest in generalizations of the two-user Gaussian interference channel model to more than 2 users and fading channels. However these generalizations turn out to be far from trivial, as they bring in new fundamental issues not encountered in the classical setting. Extensions to more than 2 users have to deal with the possibility of *interference alignment* [2], [3] while extensions to fading channels are faced with the *inseparability* of parallel interference channels [4], [5]. Interference alignment refers to the consolidation of multiple interferers into one effective entity which can be separated from the desired signal in time, frequency, space, or signal level dimensions. The inseparability of interference channels refers to the necessity for joint coding across channel states. In other words, for parallel Gaussian interference channels, the capacity cannot be expressed in general as the sum of the capacity of the sub-channels.

The following example presented in [4] to establish the inseparability of parallel interference channels forms the relevant background for this work. Consider the 3-user Gaussian interference channel with the channel matrix:

$$\mathbf{Y} = \mathbf{H}\mathbf{X} + \mathbf{Z} \tag{1}$$

where $\mathbf{Y} = [Y_1, Y_2, Y_3]^T$, $\mathbf{X} = [X_1, X_2, X_3]^T$, $\mathbf{Z} = [Z_1, Z_2, Z_3]^T$ are the vectors containing the received symbols, the transmitted symbols and the zero mean unit variance additive white Gaussian noise symbols for users indicated by the subscripts. The transmit power constraint for each user is $\mathrm{E}[X_k^2] \leq P, k = 1, 2, 3$. Consider two different values of the channel matrix,

$$\mathbf{H_a} = \begin{bmatrix} 1 & -1 & 1 \\ 1 & 1 & -1 \\ -1 & 1 & 1 \end{bmatrix}, \quad \mathbf{H}_b = \begin{bmatrix} 1 & 1 & -1 \\ -1 & 1 & 1 \\ 1 & -1 & 1 \end{bmatrix} \tag{2}$$

It is shown in [4] that taken individually either channel matrix $\mathbf{H_a}$ or $\mathbf{H_b}$ by itself results in a sum capacity of $\log(1 + 3P)$, so that separate coding can at most achieve a capacity $2\log(1 + 3P)$. However, taken together, the capacity of the parallel interference channel is $3\log(1 + 2P)$ which is achieved only by joint coding across both channel matrices. The key is the complimentary nature of the two channel matrices, i.e. $\frac{1}{2}(\mathbf{H_a} + \mathbf{H_b}) = \mathbf{I}$ which allows the receivers to cancel interference by simply adding the outputs of the parallel channels, provided the transmitters send the same symbol over both channels.

In this paper, we take this idea further by recognizing that in the ergodic setting, for a broad class of channel distributions, the channel states can be partitioned into such complimentary pairings over which interference can be aligned so that each user is able to achieve (slightly more than) half of his interference-free ergodic capacity at *any* SNR. Prior work in [3] has shown that for fading channels every user is able to achieve half the channel degrees of freedom. In other words, each user achieves (slightly less than) half of his interference-free capacity asymptotically *as SNR approaches infinity*. Fairly sophisticated interference alignment schemes are constructed to establish this achievability. However, in this work we show that for a broad class of fading distributions, including e.g. Rayleigh fading, alignment can be achieved quite simply and more efficiently. Note, however, that the stronger result is obtained at the cost of some loss of generality due to the assumption of ergodic fading and certain restrictions on the class of fading distributions, that are not needed in [3].

B. Nazer and M. Gastpar were supported by NSF grants CCR-0347298, CNS-0627024, and CCF-0830428. S. A. Jafar was supported by NSF grant CCF-0830809 and by ONR YIP under grant N00014-08-1-0872. S. Vishwanath was supported by ARO YIP under grant 52491CI.

The next section presents the main problem statement, where we formulate both a finite-field and a Gaussian interference network model. In Section III, we derive an achievable scheme for the finite field model in Section III and in Section V we show this matches the upper bound exactly. In Section IV, we give an achievable scheme for the Gaussian model which we show is quite close to the outer bound for the equal SNR case for any number of users. We conclude the paper in Section VI.

II. Problem Statement

We consider both a finite-field model and a Gaussian model. First, we will give definitions common to both models. We will use bold lowercase to denote column vectors and bold uppercase to denote matrices. There are K transmitter-receiver pairs (see Figure 1). Let n denote the number of channel uses. Let each message w_k be chosen independently and uniformly from the set $\{1, 2, \ldots, 2^{n\tilde{R}_k}\}$ for some $\tilde{R}_k \geq 0$. Message w_k is only available to transmitter k. Let $\mathcal{X}$ be the channel input and output alphabet. Each transmitter has an encoding function, $\mathcal{E}_k$, that maps the message into n channel uses:

$$\mathcal{E}_k : \{1, 2, \ldots, 2^{n\tilde{R}_k}\} \rightarrow \mathcal{X}^n \tag{3}$$

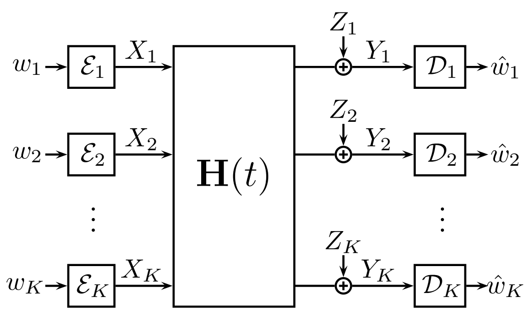

Fig. 1. K-user interference channel with fading.

We focus on the fast fading scenario where the channel matrix changes at every time step. Let $\mathbf{H}(t) = \{h_{k\ell}(t)\}_{k\ell}$ denote the channel matrix at time t and let $\mathbf{H}^n$ denote the entire sequence of channel matrices. We assume that before each time step t, all transmitters and receivers are given perfect knowledge of the channel matrix $\mathbf{H}(t)$.

At time t, the channel output seen by receiver k is given by:

$$Y_k(t) = \sum_{\ell=1}^{K} h_{k\ell}(t) X_\ell(t) + Z_k(t) \tag{4}$$

where $Z_k(t)$ is additive noise. Note that addition and multiplication are carried out over a finite field or the complex field, depending on the channel model.

Each receiver is equipped with a decoding function:

$$\mathcal{D}_k : \mathcal{X}^n \rightarrow \{1, 2, \ldots, 2^{n\tilde{R}_k}\} \tag{5}$$

and produces an estimate $\hat{w}_k$ of its desired message w_k.

Definition 1: We say that an ergodic rate tuple $(R_1, R_2, \ldots, R_K)$ is *achievable* if for all $\epsilon > 0$ and n large enough there exist channel encoding and decoding functions $\mathcal{E}_1, \ldots, \mathcal{E}_K, \mathcal{D}_1, \ldots, \mathcal{D}_K$ such that:

$$\tilde{R}_k > R_k - \epsilon, \quad k = 1, 2, \ldots, K, \tag{6}$$

$$\Pr\left(\{\hat{w}_1 \neq w_1\} \cup \ldots \cup \{\hat{w}_K \neq w_K\}\right) < \epsilon. \tag{7}$$

Definition 2: The *ergodic capacity region* is the closure of the set of all achievable ergodic rate tuples.

A. Finite Field Model

The channel alphabet is a finite field of size q, $\mathcal{X} = \mathbb{F}_q$. The channel coefficients for block n, $h_{k\ell}$, are drawn independently and uniformly from $\mathbb{F}_q \setminus \{0\}$.

Remark 1: Our results can be extended to the case where the channel coefficients are sometimes zero through simple counting arguments. However, this considerably complicates the description of the capacity region.

The additive noise terms $Z_k(t)$ are i.i.d. sequences drawn from a distribution that takes values on uniformly on $\{1, 2, \ldots, q-1\}$ with probability ρ and is zero otherwise. We define the entropy of $Z_k(t)$ to be $0 \leq H(Z) \leq \log_2 q$.

B. Gaussian Model

The channel inputs and outputs are complex numbers, $\mathcal{X} = \mathbb{C}$. Each transmitter must satisfy an average power constraint:

$$E[|X_k(t)|^2 | \mathbf{H}(t) = \mathbf{B}] \leq \mathsf{SNR}_k \quad \forall \mathbf{B} \in \mathcal{H} \tag{8}$$

where $\mathsf{SNR}_k \geq 0$ is the signal-to-noise ratio. The channel coefficients are drawn independently of each other and across time. They can be drawn from any distribution that is symmetric about zero (with $P(h_{k\ell}) = P(-h_{k\ell})$). This includes many popular fading models such as Rayleigh fading and uniform phase fading. The noise terms are i.i.d. sequences drawn from a Rayleigh distribution, $Z_k(t) \sim \mathcal{CN}(0, 1)$.

Remark 2: Our choice of power constraint eliminates the need to search for the optimal power allocation policy. A non-equal power allocation over channel states could certainly be included as part of our scheme but for the sake of simplicity we explicitly disallow it. See [6] for a study of power allocation for fast fading 2-user interference channels.

Remark 3: We could also allow for different interference-to-noise ratios between each transmitter and receiver (usually written as $\mathsf{INR}_{k\ell}$). However, the achievable rate derived in Section IV would still only depend on the SNR_k parameters.

III. Finite Field Achievable Scheme

We now develop an achievable scheme for the finite field case that can approach the symmetric ergodic capacity. First, we need some tools from the method of types [7]. Let $\mathcal{H}$ denote the alphabet of the channel matrix so that $\mathbf{H}(t) \in \mathcal{H}$. Let $N(\mathbf{H}|\mathbf{H}^n)$ be the number of times the channel matrix $\mathbf{H} \in \mathcal{H}$ occurs in the sequence $\mathbf{H}^n$.

Definition 3: A sequence of channel matrices, $\mathbf{H}^n$, is δ-typical if:

$$\left|\frac{1}{n}N(\mathbf{H}|\mathbf{H}^n) - P(\mathbf{H})\right| \leq \delta \quad \forall \mathbf{H} \in \mathcal{H} \tag{9}$$

where $P(\mathbf{H})$ is the probability of channel $\mathbf{H} \in \mathcal{H}$ under the channel model. Let A_δ^n denote the set of all δ-typical channel matrix sequences.

Lemma 1 (Csiszar-Körner 2.12): For any i.i.d. sequence of channel matrices, $\mathbf{H}^n$, the probability of the set of all δ-typical sequences, A_δ^n, is lower bounded by:

$$P(A_\delta^n) \geq 1 - \frac{|\mathcal{H}|}{4n\delta^2} \tag{10}$$

For a proof, see [7].

Lemma 2: There exists a one-to-one map, $g : \mathbb{F}_q^{K \times K} \rightarrow \mathbb{F}_q^{K \times K}$ such that $\mathbf{H} + g(\mathbf{H}) = \mathbf{I}, \quad \forall \mathbf{H}$ where $\mathbf{I}$ is the identity matrix.

Proof: Let $f : \mathbf{F}_q \rightarrow \mathbf{F}_q$ be the one-to-one map such that $f(\alpha) + \alpha = 1$ for all $\alpha \in \mathbb{F}_q$. Since $\mathbb{F}_q$ is a finite field, $f(\cdot)$ is guaranteed to exist. Then, define $g(\cdot)$ as follows:

$$g(\mathbf{H}) = \begin{bmatrix} f(h_{11}) & -h_{12} & \cdots & -h_{1K} \\ -h_{21} & f(h_{22}) & \cdots & -h_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ -h_{K1} & -h_{K2} & \cdots & f(h_{KK}) \end{bmatrix} \tag{11}$$

where $-h_{k\ell}$ is the additive inverse of $h_{k\ell}$. Clearly, $g(\mathbf{H}) + \mathbf{H} = \mathbf{I}$ and $g(\cdot)$ is one-to-one. ■

The basic idea underlying our scheme is to add together two well-chosen channel outputs such that the interference exactly cancels out. However, for the finite field model, if we do this in an uncoded fashion, we risk accumulating noise. Thus, we denoise the desired linear functions using computation codes prior to combining them together [8].

Lemma 3: Consider a K-user finite field interference channel with fixed channel coefficients $h_{k\ell} \in \mathbb{F}_q \setminus \{0\}$:

$$Y_k(t) = \sum_{\ell=1}^{K} h_{k\ell} X_\ell(t) + Z_k(t) \tag{12}$$

where $Z_k(t)$ is i.i.d. additive noise with entropy $H(Z)$. Each transmitter has a message $\mathbf{w}_k \in \mathbb{F}_q^m$. The maximum rate, $R = \frac{m}{n}\log_2 q$, at which each receiver can reliably recover the linear function $\mathbf{u}_k = \sum_{\ell=1}^{K} h_{k\ell}\mathbf{w}_\ell$ is given by:

$$R = \log_2 q - H(Z) \tag{13}$$

Proof Sketch: Let $\mathbf{G} \in \mathbb{F}_q^{n \times m}$ be a good linear code for additive noise channel at rate R. Each encoder transmits $\mathbf{x}_\ell = \mathbf{G}\mathbf{w}_\ell$. Each receiver observes:

$$\mathbf{y}_k = \sum_{\ell=1}^{K} \mathbf{G} h_{k\ell} \mathbf{w}_\ell + \mathbf{z}_k = \mathbf{G}\mathbf{u}_k + \mathbf{z}_k \tag{14}$$

from which it can recover $\mathbf{u}_k$ reliably. See Theorem 1 in [8] for a full proof and extensions.

We will now show that all users can achieve half the single user rate simultaneously.

Theorem 1: For the K-user finite field interference channel, the rate tuple $(R_{\text{SYM}}, R_{\text{SYM}}, \ldots, R_{\text{SYM}})$ is achievable where :

$$R_{\text{SYM}} = \frac{1}{2}(\log_2 q - H(Z)) \tag{15}$$

Proof: For any $\epsilon > 0$, let δ be a small positive constant that will be chosen later to satisfy our rate requirement. Using Lemma 1, choose n large enough so that $P(A_\delta^n) \geq 1 - \frac{\epsilon}{3}$. Assume that δ and n are chosen such that $n(\frac{1}{|\mathcal{H}|} - \delta)$ is an even integer. Now condition on the event that the sequence of channel matrices, $\mathbf{H}^n$, is δ-typical. Since the channel coefficients are i.i.d. and uniform, the probability of any channel $\mathbf{H} \in \mathcal{H}$ is $\frac{1}{|\mathcal{H}|}$. Since $\mathbf{H}^n$ is δ-typical we have that for every $\mathbf{H} \in \mathcal{H}$:

$$n\left(\frac{1}{|\mathcal{H}|} - \delta\right) \leq N(\mathbf{H}|\mathbf{H}^n) \leq n\left(\frac{1}{|\mathcal{H}|} - \delta\right) \tag{16}$$

Throw out all but the first $n(\frac{1}{|\mathcal{H}|} - \delta)$ indices for each channel realization. This results in losing at most a δ fraction of the total rate. Group together all time indices that have channel realization $\mathbf{H}$ and call this set of indices $\mathcal{T}_\mathbf{H}$. We will encode for each $\mathcal{T}_\mathbf{H}$ separately. For each channel realization $\mathbf{H}$, transmitter ℓ generates a message $\mathbf{w}_{\ell\mathbf{H}} \in \mathbb{F}_q^m$ where $m = \frac{n}{2}(\frac{1}{|\mathcal{H}|} - \delta)(\log_2 q)^{-1}(R_{\text{SYM}} - \frac{\epsilon}{3})$. Using a computation code from Lemma 3, each transmitter ℓ sends its message $\mathbf{w}_{\ell\mathbf{H}}$ during the first $\frac{n}{2}(\frac{1}{|\mathcal{H}|} - \delta)$ time indices in $\mathcal{T}_\mathbf{H}$. Receiver k makes an estimate $\hat{\mathbf{u}}_{k\mathbf{H}}$ of $\mathbf{u}_{k\mathbf{H}} = \sum_{\ell=1}^{K} h_{k\ell}\mathbf{w}_{\ell\mathbf{H}}$.

For each channel realization $\mathbf{H} \in \mathcal{H}$, pair up the first $\frac{n}{2}(\frac{1}{|\mathcal{H}|} - \delta)$ blocks with $\mathbf{H}$ with the last $\frac{n}{2}(\frac{1}{|\mathcal{H}|} - \delta)$ blocks with $g(\mathbf{H})$ using $g(\cdot)$ from Lemma 2. Since g is one-to-one, this procedure pairs up all of the channel indices. During the last $\frac{n}{2}(\frac{1}{|\mathcal{H}|} - \delta)$ indices with channel $g(\mathbf{H})$, the transmitters use the message, $\mathbf{w}_{\ell\mathbf{H}}$, and a computation code from Lemma 3. The receivers make an estimate $\hat{\mathbf{v}}_{k\mathbf{H}}$ of $\mathbf{v}_{k\mathbf{H}} = \mathbf{v}_{k\mathbf{H}} = f(h_{kk})\mathbf{w}_{k\mathbf{H}} - \sum_{\ell \neq k} h_{k\ell}\mathbf{w}_{\ell\mathbf{H}}$ where $f(\cdot)$ is the function such that $f(h_{k\ell}) + h_{k\ell} = 1$.

For n large enough, the total probability of error for all computation codes is upper bounded by $\frac{\epsilon}{3}$. Receiver k makes an estimate of $\mathbf{w}_{k\mathbf{H}}$ by simply adding up the two equations to get $\hat{\mathbf{w}}_{k\mathbf{H}} = \hat{\mathbf{u}}_{k\mathbf{H}} + \hat{\mathbf{v}}_{k\mathbf{H}}$. Note that the transmitters do not know a priori which time indices will be successfully paired. To deal with this, the transmitters use an erasure code with rate at least $(1-\delta)R_{\text{SYM}} - \frac{2\epsilon}{3}$ with probability of error no greater than $\frac{\epsilon}{3}$ over all transmissions. By choosing δ small enough, we finally get that each receiver can recover its message at a rate greater than $\frac{1}{2}(\log_2 q - H(Z)) - \epsilon$ with probability of error less than ϵ as desired. ■

Theorem 2: For the K-user finite field interference channel, any rate tuple $(R_1, \ldots, R_K)$, satisfying the following inequalities is achievable:

$$R_\ell + R_k \leq \log_2 q - H(Z), \quad \forall k \neq \ell. \tag{17}$$

First, we will give an equivalent description of this rate region and then show that any rate tuple can be achieved by time sharing the symmetric rate point from Theorem 1 and a single user transmission scheme.

Lemma 4: Assume, without loss of generality, that the users are labeled according to rate in descending order, so that $R_1 \geq R_2 \geq \cdots \geq R_K$. The achievable rate region from Theorem 2 is equivalent to the following rate region:

$$R_1 \leq \log_2 q - H(Z) \tag{18}$$
$$R_k \leq \min\{\log_2 q - H(Z) - R_1, \frac{1}{2}(\log_2 q - H(Z))\}, \quad k \geq 2$$

Proof: The key idea is that only one user can achieve a rate higher than $\frac{1}{2}(\log_2 q - H(Z))$. From (17), we must have that $R_1 + R_2 \leq \log_2 q - H(Z)$ so if $R_1 > \frac{1}{2}(\log_2 q - H(Z))$ all other users must satisfy $R_k \leq \log_2 q - H(Z) - R_1$. If $R_1 \leq \frac{1}{2}(\log_2 q - H(Z))$, then we have that $R_k \leq \frac{1}{2}(\log_2 q - H(Z))$ for all other users since the rates are in descending order. ■

Proof of Theorem 2: We show that the equivalent rate region developed by Lemma 4 is achievable by time-sharing. First, we consider the case where $R_1 > \frac{1}{2}(\log_2 q - H(Z))$. Let $\alpha = 2(1 - \frac{R_1}{\log_2 q - H(Z)})$. We allocate αn channel uses to the symmetric scheme from Theorem 1. For, the remaining $(1-\alpha)n$ channel uses, users 2 through K are silent, and user 1 employs a capacity-achieving point-to-point channel code. This results in user 1 achieving its target rate R_1:

$$\begin{aligned}&\frac{\alpha(\log_2 q - H(Z))}{2} + (1-\alpha)(\log_2 q - H(Z)) \qquad (19)\\ &= \log_2 q - H(Z) - R_1 - \log_2 q + H(Z) + 2R_1 = R_1\end{aligned}$$

and users 2 through K achieving $R_k = \log_2 q - H(Z) - R_1$. If $R_1 \leq \frac{1}{2}(\log_2 q - H(Z))$, we can achieve any rate point with the use of the symmetric scheme from Theorem 1. ■

IV. Gaussian Achievable Scheme

The scheme for the Gaussian case is quite similar to our finite field scheme. The key difference is that we need to quantize the channel alphabet so that we can deal with a finite set of possible matrices. By decreasing the quantization bin size, we can approach the desired rate in the limit. Also, here it is beneficial to transmit combine the channel outputs prior to decoding to exploit a power gain.

Definition 4: For $\gamma > 0$, let $Q_\gamma(h_{k\ell})$ represent the closest point in $\gamma(\mathbb{Z} + j\mathbb{Z})$ to $h_{k\ell}$ in Euclidean distance. The γ-*quantized* version of a channel matrix $\mathbf{H} \in \mathbb{C}^{K \times K}$ is given by $\mathbf{H}_\gamma = \{Q_\gamma(h_{k\ell})\}_{k\ell}$.

Theorem 3: For the K-user Gaussian interference channel, the rate tuple $(R_1, R_2, \ldots, R_K)$ is achievable for :

$$R_k = \frac{1}{2} E\left[\log\left(1 + 2|h_{kk}|^2 \mathsf{SNR}_k\right)\right]. \tag{20}$$

Proof: For any $\epsilon > 0$, choose $\tau > 0$ such that $P\left(\cup_{k\ell}\{|h_{k\ell}| > \tau\}\right) < \frac{\epsilon}{3}$. Let γ and δ be small positive constants that will be chosen later to satisfy our rate requirement. Also, using Lemma 1, choose n large enough so that $P(A_\delta^n) \geq 1 - \frac{\epsilon}{3}$. We will throw out any time index with a channel coefficient with magnitude larger than τ. This ensures that the γ-quantized version of the channel is of finite size. Specifically, the size of the channel alphabet $\mathcal{H}$, is given by $|\mathcal{H}| = (2\frac{\tau}{\gamma})^{2K^2}$. We assume that τ, γ, δ and n are chosen so that all the appropriate ratios only result in integers.

We condition on the event that the sequence of γ-quantized channel matrices, $\mathbf{H}_\gamma^n$, is δ-typical. Unlike the finite field case, the channel matrix distribution is not uniform. For all $\mathbf{H}_\gamma \in \mathcal{H}$ we have that:

$$n(P(\mathbf{H}_\gamma) - \delta) \leq N(\mathbf{H}_\gamma | \mathbf{H}_\gamma^n) \leq n(P(\mathbf{H}_\gamma) + \delta) \tag{21}$$

Throw out all but the first $n(P(\mathbf{H}_\gamma) - \delta)$ blocks of each channel realization. This causes a loss of at most a δ fraction in rate. Let $h_{k\ell}^\gamma$ denote the elements of $\mathbf{H}_\gamma$. We define the following one-to-one map $g : \mathcal{H} \rightarrow \mathcal{H}$:

$$g(\mathbf{H}_\gamma) = \begin{bmatrix} h_{11}^\gamma & -h_{12}^\gamma & \cdots & -h_{1K}^\gamma \\ -h_{21}^\gamma & h_{22}^\gamma & \cdots & -h_{2K}^\gamma \\ \vdots & \vdots & \ddots & \vdots \\ -h_{K1}^\gamma & -h_{K2}^\gamma & \cdots & h_{KK}^\gamma \end{bmatrix} \tag{22}$$

Note that due to the symmetry of the channel distribution $P(g(\mathbf{H}_\gamma)) = P(\mathbf{H}_\gamma)$. Group together all time indices that have channel realization $\mathbf{H}_\gamma$ and call this set of indices $\mathcal{T}_{\mathbf{H}_\gamma}$. For each channel realization $\mathbf{H} \in \mathcal{H}$, pair up the first $\frac{n}{2}(P(\mathbf{H}_\gamma) - \delta)$ blocks with channel $\mathbf{H}_\gamma$ with the last $\frac{n}{2}(P(\mathbf{H}_\gamma) - \delta)$ blocks with channel $g(\mathbf{H}_\gamma)$. We ensure that we use the same channel inputs during time index i from $\mathcal{T}_{\mathbf{H}_\gamma}$ for $i = 1, 2, \ldots, \frac{n}{2}(P(\mathbf{H}_\gamma) - \delta)$ as we do during time index $i + \frac{n}{2}(P(\mathbf{H}_\gamma) - \delta)$ from $\mathcal{T}_{g(\mathbf{H}_\gamma)}$. Let t_1 denote the first time and t_2 denote the second time. We have the following channel outputs:

$$Y_k(t_1) = h_{kk}(t_1)X_k(t_1) + \sum_{\ell \neq k} h_{k\ell}(t_1)X_\ell(t_1) + Z_k(t_1)$$
$$Y_k(t_2) = h_{kk}(t_2)X_k(t_1) + \sum_{\ell \neq k} h_{k\ell}(t_2)X_\ell(t_1) + Z_k(t_2)$$

Since t_1 has quantized channel $\mathbf{H}_\gamma$ and t_2 has quantized channel $g(\mathbf{H}_\gamma)$ we have that the channel from $X_k(t_1)$ to $Y_k(t_1) + Y_k(t_2)$ has a signal-to-noise ratio of at least:

$$\frac{\mathsf{SNR}_k(2(Re(h_{kk}) - \frac{\gamma}{2})^2 + (Im(h_{kk}) - \frac{\gamma}{2})^2)}{2 + \gamma^2 \sum_{\ell \neq k} \mathsf{SNR}_\ell} \tag{23}$$

By choosing γ small enough, we can achieve:

$$R_{k\mathbf{H}_\gamma} > \max_{h_{kk} \in \mathbf{H}_\gamma} \frac{1}{2} \log\left(1 + 2|h_{kk}|^2 \mathsf{SNR}_k\right) - \frac{\epsilon}{3} \tag{24}$$

for each $\mathbf{H}_\gamma$. The total rate per user is given by

$$R_k = \frac{1}{|\mathcal{H}|} \sum_{\mathbf{H}_\gamma \in \mathcal{H}} P(\mathbf{H}_\gamma) R_{k\mathbf{H}_\gamma}(1 - \delta) \tag{25}$$

For δ small enough and taking the limit $\gamma \rightarrow 0$, we get:

$$\lim_{\gamma \rightarrow 0} R_k =$$
$$\frac{1}{2} \int 1\{|h_{k\ell}| > \tau\} \log\left(1 + 2|h_{kk}|^2 \mathsf{SNR}_k\right) P(\mathbf{H}) d\mathbf{H} - \frac{2\epsilon}{3}$$

Finally, taking $\tau \rightarrow \infty$, we get:

$$\lim_{\tau \rightarrow \infty} \lim_{\gamma \rightarrow 0} R_k = \frac{1}{2} E[\log\left(1 + 2|h_{kk}|^2 \mathsf{SNR}_k\right)] - \frac{2\epsilon}{3} \tag{26}$$

Thus, there exist γ and τ such that we achieve $R_k > \frac{1}{2} E[\log\left(1 + 2|h_{kk}|^2 \mathsf{SNR}_k\right)] - \epsilon$ with probability $1 - \epsilon$. ■

V. Upper Bounds

We now briefly describe upper bounds for both the finite field case and the Gaussian case. The finite field upper bound matches the achievable performance thus yielding the ergodic capacity region. For the Gaussian case, we demonstrate that our achievable performance is very close to the upper bound when the transmitters have equal power constraints.

Theorem 4: For the K-user finite field interference channel, the ergodic capacity region is:

$$R_\ell + R_k \leq \log_2 q - H(Z), \qquad \forall k \neq \ell. \tag{27}$$

Proof: The required upper bound follows from steps similar to those in Appendix II of [3]. Without loss of generality, we upper bound the rates of users 1 and 2. Note that the capacity of the interference channel only depends on the noise marginals. Thus, we can assume that $Z_1(t) = h_{12}(t)(h_{22}(t))^{-1}Z_2(t)$. Let $\tilde{Y}_2(t) = h_{12}(t)(h_{22}(t))^{-1}Y_2(t)$.

We give the receivers full access to the messages from users 3 through K as this can only increase the outerbound. From Fano's inequality, we have that $n(R_1+R_2-\epsilon_n)$ where $\frac{\epsilon_n}{n} \to 0$ as $n \to \infty$ is upper bounded as follows:

$$\begin{aligned}
&\leq I(w_1;Y_1^n) + I(w_2;w_1,\tilde{Y}_2^n) \\
&= I(w_1;Y_1^n) + I(w_2;\tilde{Y}_2^n|w_1,X_1^n) \\
&= I(w_1;Y_1^n) + I(w_2;\{h_{12}(t)X_2(t)+Z_1(t)\}_{t=1}^n|w_1,X_1^n) \\
&= I(w_1;Y_1^n) + \\
&\quad\; I(w_2;\{h_{11}(t)X_1(t)+h_{12}(t)X_2(t)+Z_1(t)\}_{t=1}^n|w_1,X_1^n) \\
&= I(w_1;Y_1^n) + I(w_2;Y_1^n|w_1) \\
&= I(w_1,w_2;Y_1^n) \\
&\leq n(\log_2 q - H(Z))
\end{aligned}$$

Similar outer bounds hold for all receiver pairs k and ℓ. Comparing these to the achievable region in Theorem 2 yields the capacity region. ∎

Using the results from [6], we have the following outer bound on the ergodic capacity region of the K-user Gaussian interference channel.

Theorem 5: For the K-user Gaussian interference channel with i.i.d. Rayleigh fading, the following constraints are an outer bound to the ergodic capacity region:

$$\begin{aligned}
R_k + R_\ell \leq\ & E\left[\log\left(1 + |h_{k\ell}|^2\mathsf{SNR}_\ell + \frac{|h_{kk}|^2\mathsf{SNR}_k}{1+|h_{\ell k}|^2\mathsf{SNR}_k}\right)\right] \\
+\ & E\left[\log\left(1 + |h_{\ell k}|^2\mathsf{SNR}_k + \frac{|h_{\ell\ell}|^2\mathsf{SNR}_\ell}{1+|h_{k\ell}|^2\mathsf{SNR}_\ell}\right)\right] \forall k \neq \ell
\end{aligned}$$

In Figure 2, we plot the performance of our scheme versus the upper bound from Theorem 5 for the equal SNR, equal rate per user case. The plot is for i.i.d. Rayleigh fading and is valid for any number of users K. This shows that ergodic interference alignment can provide close-to-optimal performance for any number of users so long as they have the same SNR constraint. In very recent work, Jafar has shown that ergodic interference alignment achieves capacity whenever a network is in a "bottleneck state." This includes, as a special case, uniform phase fading channels with a large number of users [9].

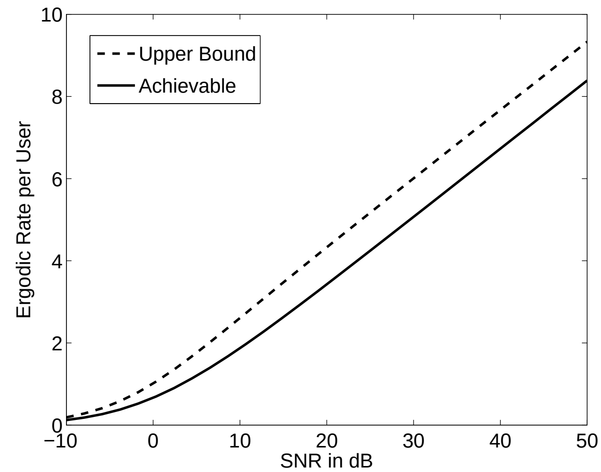

Fig. 2. Ergodic symmetric rate per user and upper bound for the K-user Gaussian interference channel with i.i.d. Rayleigh fading and equal transmit powers: $\mathsf{SNR}_1 = \mathsf{SNR}_2 = \cdots = \mathsf{SNR}_K$.

VI. Conclusions

We developed a new communication strategy, ergodic interference alignment, that codes efficiently across parallel interference channels. With this strategy, every user in the channel can attain at least half the rate available to them in the single-user setting. Moreover, we showed that for a finite field channel model this achievable scheme matches the outer bound exactly, thus yielding the ergodic capacity region. The key to the achievable strategy was perfect channel knowledge at the transmitters. An interesting direction for future work is developing an ergodic alignment scheme for the case of limited channel state information.

References

[1] R. H. Etkin, D. N. C. Tse, and H. Wang, "Gaussian interference channel capacity to within one bit," *IEEE Transactions on Information Theory*, vol. 54, pp. 5534–5562, December 2008.

[2] M. A. Maddah-Ali, A. S. Motahari, and A. K. Khandani, "Communication over mimo x channels: Interference alignment, decomposition, and performance analysis," *IEEE Transactions on Information Theory*, vol. 54, pp. 3457–3470, August 2008.

[3] V. R. Cadambe and S. A. Jafar, "Interference alignment and the degrees of freedom for the K user interference channel," *IEEE Transactions on Information Theory*, vol. 54, pp. 3425–3441, August 2008.

[4] V. Cadambe and S. A. Jafar, "Multiple access outerbounds and the inseparability of parallel interference channels," Dec. 2008.

[5] L. Sankar, X. Shang, E. Erkip, and H. V. Poor, "Ergodic two-user interference channels: Is separability optimal?," *46th Annual Allerton Conf. on Comm., Control, and Computing*, 2008.

[6] D. Tuninetti, "Gaussian fading interference channels: Power control," in *Proceedings of the 42nd Asilomar Conference on Signals, Systems and Computers*, (Monterey, CA), October 2008.

[7] I. Csiszár and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*. New York: Academic Press, 1982.

[8] B. Nazer and M. Gastpar, "Computation over multiple-access channels," *IEEE Transactions on Information Theory*, vol. 53, pp. 3498–3516, October 2007.

[9] S. A. Jafar, "The ergodic capacity of interference networks," tech. rep., October 2009. See http://arxiv.org/abs/0902.0838.