

# Interference Alignment — A New Look at Signal Dimensions in a Communication Network

By Syed A. Jafar

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Linear Interference Alignment — Concept</b>	<b>4</b>
2.1	Few Observed Equations, Many Unknowns	6
<b>3</b>	<b>Origins of Interference Alignment</b>	<b>10</b>
3.1	Index Coding	10
3.2	X Channel	12
3.3	Interference Channel with $K > 2$ Users	14
<b>4</b>	<b>New Challenges and Solutions</b>	<b>18</b>
4.1	Feasibility of Linear Interference Alignment	19
4.2	Symbol Extensions	26
4.3	Asymmetric Complex Signaling	27
4.4	Channel Variations: Separability	33
4.5	Ergodic Interference Alignment	37

4.6	An Asymptotic Interference Alignment Scheme — [CJ08]	39
4.7	Interference Alignment Based on Separability of Rationally Independent Dimensions	49
4.8	Lattice Alignment	52
4.9	Blind Interference Alignment	59
4.10	Retrospective Interference Alignment	63
<b>5</b>	<b>Applications of Interference Alignment</b>	<b>68</b>
5.1	$K$ User Interference Channel	69
5.2	$K$ User $M \times N$ MIMO Interference Channels	71
5.3	Cellular Networks	75
5.4	X Networks	79
5.5	Compound MISO BC Channel	81
5.6	Network Coding — Multiple Unicasts	83
5.7	Distributed Storage Exact Repair Problem	86
5.8	Multihop Interference Networks	92
5.9	Bidirectional Relay Interference Networks	99
5.10	Cooperative Interference Networks	102
5.11	Secrecy	106
<b>6</b>	<b>Conclusion</b>	<b>109</b>
<b>A</b>	<b>Degrees of Freedom (DoF)</b>	<b>112</b>
<b>B</b>	<b>Generalized Degrees of Freedom (GDoF)</b>	<b>115</b>
	<b>References</b>	<b>122</b>

Foundations and Trends<sup>®</sup> in  
Communications and Information Theory  
Vol. 7, No. 1 (2010) 1–134  
© 2011 S. A. Jafar  
DOI: 10.1561/01000000047



# Interference Alignment — A New Look at Signal Dimensions in a Communication Network

Syed A. Jafar

*University of California, Irvine, CA 92697, USA, syed@uci.edu*

## **Abstract**

This monograph introduces to the reader the idea of interference alignment, traces its origins, reviews a variety of interference alignment schemes, summarizes the diverse settings where the idea of interference alignment is applicable and highlights the common principles that cut across these diverse applications. The focus is on theoretical aspects.

# 1

---

## Introduction

---

Interference alignment is a radical idea that has recently emerged out of the capacity analysis of interference networks. In a relatively short time, this concept has challenged much of the conventional wisdom about the throughput limits of both wired and wireless networks. A representative example is the wireless interference channel with  $K$  transmitter–receiver pairs where, because of interference alignment, each user is simultaneously able to send at a data rate equal to half of his interference-free channel capacity to his desired receiver, even though the number of users  $K$  can be arbitrarily large, thus showing that the interference channel is not fundamentally interference-limited. While the remarkable benefits of interference alignment have so far been shown mostly under idealized assumptions such as global channel knowledge, bandwidth expansion, unlimited resolution, high signal strengths and significant delays, the idea has garnered rapidly increasing interest in the communication, signal processing, networking and information theory communities and has produced an array of surprising and fundamental insights into the number of accessible signaling dimensions in both wired and wireless communication networks. A diversity of tools from linear algebra, algebraic geometry, diophantine approximation theory as well as coding and

traditional Shannon theory continue to be the basis for an increasing variety of interference alignment schemes that include spatial alignment, lattice alignment, asymptotic alignment, asymmetric complex signal alignment, opportunistic alignment, ergodic alignment, aligned interference neutralization, blind alignment and retrospective alignment schemes. Applications include wireless interference networks, X networks, cellular networks, two-way communication networks, multicast and compound networks, multihop multiflow networks, tactical communication networks with secrecy and jamming issues, cooperative communication networks, cognitive radio networks, distributed data storage networks, index coding networks, and wired multiple unicast networks.

The goal of this monograph is to provide both a tutorial and a survey of the state-of-the-art on the topic of interference alignment. The majority of the paper is written to be accessible to a graduate student working in communication, signal processing, networking or information theory. The focus is on theoretical aspects. The presentation style of this monograph is informal, favoring broad intuition over mathematical rigor, and the finer details are mostly left to the references.

This monograph is organized into six sections, starting with this section — the introduction. Section 2 introduces the fundamental idea of interference alignment in its simplest form — in the language of elementary linear algebra. The origins of interference alignment are the topic of Section 3 where the earliest applications of interference alignment are reviewed. Section 4 explores the challenges faced by interference alignment schemes and the ingenious solutions that get around those challenges. Section 5 reviews the diverse settings, including both wireless and wired networks, where interference alignment has been applied while highlighting how many of these diverse problems can be reduced to the same essential problem for which a systematic solution already exists. Each subsection of this section ends with pointers to one or more open problems. Section 6 concludes the paper. An Appendix is included at the end which provides a rudimentary introduction to the Degrees-of-Freedom (DoF) and Generalized Degrees of Freedom (GDoF) metrics that are necessary to appreciate the concepts discussed throughout the paper.

# 2

---

## Linear Interference Alignment — Concept

---

While there are many interference alignment schemes that take much more sophisticated forms, the origins of the idea lie in elementary linear algebra. Consider a system of linear equations

$$\begin{aligned}y_1 &= h_{11}x_1 + h_{12}x_2 + \cdots + h_{1K}x_K \\y_2 &= h_{21}x_1 + h_{22}x_2 + \cdots + h_{2K}x_K \\&\vdots \\y_B &= h_{B1}x_1 + h_{B2}x_2 + \cdots + h_{BK}x_K,\end{aligned}\tag{2.1}$$

where we have  $B$  observations  $y_1, y_2, \dots, y_B$ , each in the form of a linear combination of  $K$  information symbols  $x_1, x_2, \dots, x_K$  with coefficients  $h_{ij}$ . Since we are interested in interference networks, let us interpret  $K$  to be the number of transmitters, each trying to send one information symbol, for a total of  $K$  independent information symbols  $x_1, x_2, \dots, x_K$ . The coefficients  $h_{ij}$  may be interpreted as the effective *channel* coefficients. Further, let us interpret  $B$  to be the bandwidth, or the number of signaling dimensions accessible at a receiver through a linear channel. Since the channel is linear, each signaling dimension

produces a linear combination of the transmitted information symbols. Thus, a receiver has access to  $B$  signaling dimensions.<sup>1</sup>

If the equations are generic, e.g., if effective channel coefficients are drawn from a continuous distribution, all the information symbols can be recovered, provided that there are at least as many observations as unknowns. Thus, if all symbols are desired by the receiver, we need at least  $K$  signaling dimensions. However, we are interested in *interference* networks where only a subset of the symbols are desired by the receiver and the remaining symbols (which carry information for *other* receivers) are undesired at this receiver, i.e., they only contribute interference. As a specific example, suppose the receiver is only interested in symbol  $x_1$  which carries information desired by this receiver while all other symbols carry only interference. In this setting, we ask the question — how many observations does this receiver need in order to be able to resolve its desired symbol from the interference? In other words, how many signaling dimensions are needed, or equivalently, how much bandwidth is needed for the receiver to be able to resolve its desired symbol from the remaining  $K - 1$  undesired symbols?

In general,  $K$  signaling dimensions will be needed to resolve the 1 symbol desired by this receiver. And since there are presumably  $K$  receivers, each interested in a different symbol, and each with access to a different set of  $K$  linear equations dictated by its linear channel to the transmitters, each receiver will be able to solve the system of equations and recover its desired symbol. Thus, a total of  $K$  signaling dimensions (or a bandwidth of  $K$ ) are used so that each receiver is able to resolve its desired one dimensional signal. In the terminology of wireless interference networks, this solution corresponds to the cake-cutting interpretation of spectrum allocation — the total number of signaling dimensions, i.e., the total bandwidth, is divided among the  $K$

---

<sup>1</sup>To keep the discussion general as well as compact we do not specify if the symbols are from a finite field or if they are real/complex, as the same ideas apply in both cases. In the latter case, however, it is worth mentioning that as usual there are additive noise terms and power constraints present, that are not shown in these equations, but without which the communication problem would be trivial, i.e., the infinite resolution of a real/complex number will allow infinite information transfer in any signal dimension.

users much like a cake, so that each user can communicate over a  $1/K$  fraction of it.

However, as it turns out, the cake-cutting bandwidth allocation is not optimal. It is possible to recover the desired information symbol even when the number of linear equations available to the receiver is much smaller than the number of unknowns.

## 2.1 Few Observed Equations, Many Unknowns

To understand this, let us ask the question — what is required for a receiver to resolve symbol  $x_1$  from interference? This question has a simple answer. Let us re-write the system of linear equations observed by the receiver as follows:

$$\mathbf{Y} = \mathbf{H}_{*1}x_1 + \mathbf{H}_{*2}x_2 + \cdots + \mathbf{H}_{*K}x_K, \quad (2.2)$$

where

$$\mathbf{Y} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_B \end{bmatrix} \quad \text{and} \quad \mathbf{H}_{*k} = \begin{bmatrix} h_{1k} \\ h_{2k} \\ \vdots \\ h_{Bk} \end{bmatrix}$$

are the observations vector and the signaling dimension along which the symbol  $x_k$  is observed, respectively. In the parlance of the multiple-input multiple-output systems (MIMO) literature, we will call the vector  $\mathbf{H}_{*k}$  the received *beam* direction for symbol  $x_k$ . With this terminology, the condition for the receiver to recover the desired symbol  $x_1$  from the observations vector  $\mathbf{Y}$  is simply that the received *beam* for  $x_1$ , i.e.,  $\mathbf{H}_{*1}$  should not be contained in the vector space spanned by the undesired beams  $\mathbf{H}_{*2}, \dots, \mathbf{H}_{*K}$ . In other words, *the receiver can recover  $x_1$  if and only if*

$$\mathbf{H}_{*1} \notin \text{span}(\mathbf{H}_{*2}, \mathbf{H}_{*3}, \dots, \mathbf{H}_{*K}). \quad (2.3)$$

The  $K - 1$  interference beams  $\mathbf{H}_{*2}, \mathbf{H}_{*3}, \dots, \mathbf{H}_{*K}$  will *in general* span a vector space of  $\min(B, K - 1)$  dimensions. Thus, if  $B < K$ , i.e., the bandwidth is smaller than the number of users, the interference will span  $\min(B, K - 1) = B$  dimensions. Since all the available  $B$  dimensions are spanned by interference beams, the desired signal beam will lie

within the interference space as well and cannot be resolved. However, if the interference beams could be consolidated into a smaller subspace so that they do not span the entire available signal space at the receiver, and the desired signal beam could avoid falling into the interference space, then the receiver could indeed recover its desired symbol. This is precisely the idea of interference alignment.

For example, consider the system of observed equations

$$y_1 = 3x_1 + 2x_2 + 3x_3 + x_4 + 5x_5 \quad (2.4)$$

$$y_2 = 2x_1 + 4x_2 + x_3 - 3x_4 + 5x_5 \quad (2.5)$$

$$y_3 = 4x_1 + 3x_2 + 5x_3 + 2x_4 + 8x_5. \quad (2.6)$$

So we are operating over a three-dimensional space (total bandwidth = 3), and there are five unknowns. To solve for all five unknowns we will need five observations (i.e., total bandwidth = 5). However, the receiver only wants to recover  $x_1$ . From these equations, even though the number of equations is smaller than the number of unknowns, it turns out that indeed  $x_1$  can be recovered, because the interfering beams span only a two-dimensional vector space, leaving one dimension free from interference. To see this, note that  $\mathbf{H}_{*4} = \mathbf{H}_{*3} - \mathbf{H}_{*2}$  and  $\mathbf{H}_{*5} = \mathbf{H}_{*3} + \mathbf{H}_{*2}$ . The received signal can be projected along the interference free dimension to recover  $x_1$ . To see this explicitly, note that the vector  $\mathbf{U} = [17 \ -1 \ -10]^T$  is orthogonal to all the interference vectors, and therefore the projection  $\mathbf{U}^T \mathbf{Y}$  eliminates all the interference leaving us:

$$17y_1 - y_2 - 10y_3 = 9x_1 \quad (2.7)$$

thus,  $x_1$  is recovered from the observed values  $y_1, y_2, y_3$  in the three-dimensional space, even when the number of unknowns is  $5 > 3$ .

To summarize, interference alignment allows many interfering users to communicate simultaneously over a small number of signaling dimensions (bandwidth) by consolidating the space spanned by the interference at each receiver within a small number of dimensions, while keeping the desired signals separable from interference so that they can be projected into the null space of the interference and thereby recovered *free from interference*.

The desired outcome sought by interference alignment is illustrated by the simple example provided above. How to accomplish this goal within the constraints of each specific communication network is the main challenge in designing linear interference alignment schemes. The example presented above focused on one receiver (say, Receiver 1) that desires the symbol  $x_1$ . Presumably, there are other receivers simultaneously observing the outputs of their own channels over the same bandwidth  $B$ , i.e., they each have  $B$  linear equations, but are interested in different desired symbols. For instance, another receiver (say Receiver 2) may want only the symbol  $x_2$  and the symbols  $x_1, x_3, x_4, x_5$  will constitute interference at that receiver. For Receiver 2 to be able to recover  $x_2$  he needs the interfering beams due to symbols  $x_1, x_3, x_4, x_5$  to align into a two dimensional space and the desired signal beam for  $x_2$  to not align in that space. The reason the requirements of Receiver 2 do not directly conflict with the requirements of Receiver 1 is because each receiver sees a different picture, i.e., the set of equations seen at each receiver is different from other receivers. This is the very important principle of *relativity of alignment*, i.e., the alignment of signals is different at each receiver. Since each receiver sees a different picture, it may be possible through careful design of transmitted signals, to tailor the observations of each receiver in such a manner that the desired signals remain resolvable, while the interfering signals cast *overlapping shadows*. Thus, the relativity of alignment is the enabling premise of interference alignment. Without it, the opportunity for interference alignment would not exist.

While the relativity of alignment opens the door to interference alignment schemes, it does not by itself solve the problem for us. Signals do not align into desirable patterns naturally, or by accident. Much like when dropping a couple of pencils on the floor, the probability that they will land exactly parallel is small, so is the probability that a random choice of signal vectors will automatically achieve desired interference alignment outcomes by accident. Thus, the most important challenge in interference alignment schemes is the design of signal vectors to satisfy the desired alignment conditions. Different settings allow varying amounts of freedom in this regard. For example, the channel coefficients in one-hop wireless networks are typically determined

entirely by nature, and interference must be aligned in spite of having no control over the channel realizations. In multihop wired networks, e.g., network coding problems, or even in multihop wireless networks, the choice of the functionality of intermediate nodes may provide some amount of additional freedom to design the end-to-end channel itself to make alignment easier. In most cases, however, the alignment problem remains nontrivial.

# 3

---

## Origins of Interference Alignment

---

Applications of interference alignment can be found as early as the 1998 paper of Birk and Kol [13, 14] where the index coding problem was introduced. Interference alignment was observed again in 2006 in the specialized context of the X channel by Maddah-Ali et al. in [90] and was implicitly used for the compound MISO broadcast (BC) channel by Weingarten et al. in [159] in 2007. The idea was crystallized by Jafar and Shamai in [67] and its strength as a general principle was established by Cadambe and Jafar who introduced in [22] a mechanism to align an arbitrarily large number of interferers, leading to the surprising conclusion that wireless networks are not essentially interference limited. It has since continued to evolve in increasingly sophisticated forms across a variety of applications. In this section, we review the early applications of interference alignment in [13, 14, 22, 67, 90].

### 3.1 Index Coding

The earliest application of interference alignment, that we are aware of, appears in Example 7 of the 1998 INFOCOM paper of Birk and Kol [13, 14] in the context of the Informed Source Coding on Demand

(ISCOD) problem (also known as the index coding problem). The problem formulation, introduced in the same paper, can correspond to a variety of settings. As described by Birk and Kol, it can be seen as a wireless BC channel with cognitive receivers. It can also equivalently be formulated as a network coding problem over wired networks.

Figure 3.1 presents the problem defined in Example 7 of Birk and Kol’s paper both as a wireless BC channel with cognitive receivers as well as the equivalent wired network coding problem. This is a BC channel setting where five independent information symbols  $a, b, c, d, x$  are desired by five separate receivers (labeled according to their desired messages in Figure 3.1). Each receiver has some cognitive side information, i.e., knows some of the messages that it does not desire (e.g., from prior transmissions of those messages that failed to reach the desired nodes but were strong enough to be decoded at the undesired receivers). For instance, in Figure 3.1, information symbols  $b, x$  are known at the receiver that desires symbol  $a$ ;  $a, x$  are known by the receiver that desires  $b$ ;  $b, d$  are known by the receiver that desires  $c$ ;  $b, c$  are known at the receiver that desires  $d$ ; and  $a, c, d$  are known at the receiver that

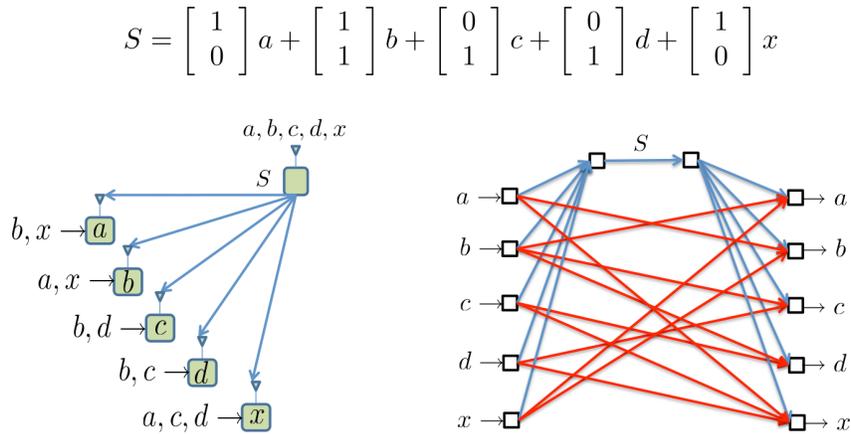


Fig. 3.1 An interference alignment solution — Example 7 of [13]. On the left is a BC channel with cognitive receivers. On the right is a network coding version of the same problem. In each case, after removing the two known interference symbols available to them, the receivers that desire  $a, b, c, d$ , respectively, each have three unknowns left and only two linear equations provided by  $S$ . In each case interference alignment is employed to align the two remaining undesired symbols into one dimension, leaving the other dimension interference-free to recover the desired symbol.

desires the symbol  $x$ . The solution proposed by Birk and Kol uses two signaling dimensions, i.e., two consecutive channel uses, to send the symbol  $S$  from the transmitter which is comprised as:

$$S = \begin{bmatrix} 1 \\ 0 \end{bmatrix} a + \begin{bmatrix} 1 \\ 1 \end{bmatrix} b + \begin{bmatrix} 0 \\ 1 \end{bmatrix} c + \begin{bmatrix} 0 \\ 1 \end{bmatrix} d + \begin{bmatrix} 1 \\ 0 \end{bmatrix} x. \quad (3.1)$$

To see how interference alignment is used, let us consider how each receiver recovers its desired information. Receiver  $a$ , i.e., the receiver that desires the symbol  $a$  has access to the two-dimensional symbol  $S$ , and it also observes  $b$  and  $x$  as cognitive information. After removing the known symbols  $b$  and  $x$  from  $S$  the receiver is left with three remaining unknown symbols  $a, c, d$  in a two-dimensional space. Since  $a$  is the desired symbol, the remaining two symbols  $c, d$  that constitute interference, must align. Indeed, since both  $c$  and  $d$  are seen along the interference beam  $[0 \ 1]^T$  these symbols do align into a one-dimensional space, and since  $a$  is seen along the beam  $[1 \ 0]^T$ , it remains resolvable from the interference. Similarly, at receiver  $b$ , after removing the known symbols there are three remaining unknowns  $b, c, d$  in a 2 dimensional space. Once again, since  $c, d$  align and  $b$  does not align with  $c, d$  the desired information is recovered. Now consider receiver  $c$  which is left with three unknown symbols  $a, c, x$  in a two-dimensional space. However, since  $a, x$  are both aligned into the same dimension  $[1 \ 0]^T$ , and  $c$  is along a different dimension, the desired information is recovered. Interference alignment is similarly exploited at Receiver  $d$ . And finally, receiver  $x$ , is left with only two unknown symbols  $b, x$  in a two-dimensional space. Since two equations are enough to recover two unknown symbols, no interference alignment is needed at Receiver  $x$ .

Birk and Kol point out the surprising nature of this result, and attribute the “apparent miracle” to the observation that not all information is desired by all nodes, which is indeed the key to interference alignment, allowing undesired information to be consolidated into a smaller dimension.

### 3.2 X Channel

While the index coding literature continued to use interference alignment schemes specialized to their problem, the study of the DoF of

the two user wireless X channel initiated by Maddah-Ali et al. in a 2006 ISIT paper [91] found surprisingly high DoF (relative to the prior studies of the interference channel [65]). The scheme proposed by Maddah-Ali et al. was not an interference alignment scheme per se — it was an iterative achievable scheme for the X channel, built upon dirty paper coding and successive decoding as integral elements, that was observed by Maddah-Ali et al. in [90] to lead to “increasing overlap of interference spaces” as signal-to-noise power ratio (SNR) approaches infinity.

The surprising DoF results of [90, 91] and the observation of the role of overlapping interference spaces motivated the work of Jafar et al., starting as a 2006 technical report [62] and leading up to the results in [67]. The most significant contribution of Jafar and Shamai’s work was to crystallize the idea of interference alignment. Specifically, reference [67] presented the interference alignment scheme for the two user X channel in its essential linear form (see Figure 3.2), thereby establishing that neither dirty paper coding nor successive decoding or iterative solutions were required. By stripping away the restrictive dependencies on the X channel setting and presenting alignment from a linear algebraic vector space perspective, the explicit alignment solution of [67] showed the idea of interference alignment to be a *general* principle. The recognition of the generality of the idea also necessitated a context-independent nomenclature — and the terminology *interference alignment* was introduced [67]. Following [62, 67], an explicit linear form of interference alignment was subsequently also adopted by Maddah-Ali et al. in their full paper [89, 92].

Figure 3.2 shows an X channel with two transmitters and two receivers, and all nodes equipped with three antennas. The channel between Transmitter  $t$  and Receiver  $r$  is assumed to be a generic<sup>1</sup>  $3 \times 3$  linear transformation indicated as  $\mathbf{H}^{[rt]}$  with  $r, t \in \{1, 2\}$ . Each transmitter has a message for each receiver for a total of four independent messages. Even though there are a total of four messages and only three signaling dimensions interference alignment enables one symbol

---

<sup>1</sup>The “generic” characterization refers to the assumption that the channel coefficients are drawn from a continuous distribution.

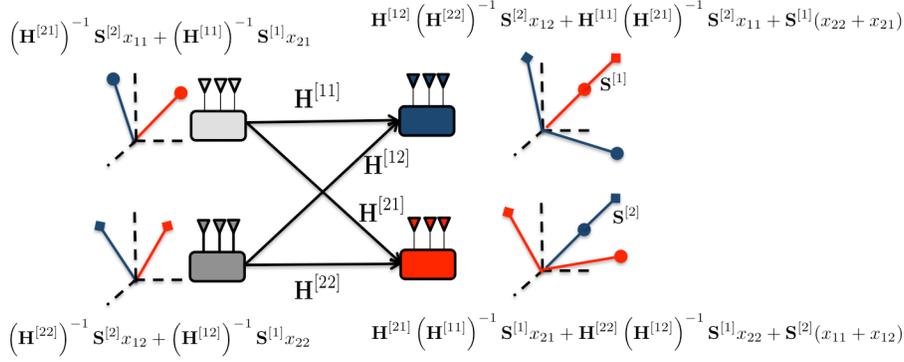


Fig. 3.2 Interference alignment on the X Channel [67] — Each receiver sees three linear equations in four unknowns. Since only two of the unknowns are desired Interference alignment makes it possible to recover their values by aligning the two undesired symbols along the same dimension at each receiver.

per message to be delivered free from interference. This is accomplished by aligning the two undesired symbols at each receiver along a one-dimensional vector, indicated as the generic  $3 \times 1$  vector  $\mathbf{S}^{[1]}$  for Receiver 1 and  $\mathbf{S}^{[2]}$  for Receiver 2 in Figure 3.2. The choice of beamforming vectors at the transmitter is such that the received vectors will be aligned in the direction  $\mathbf{S}^{[i]}$  for the undesired receiver  $i$ . For instance, the symbol  $x_{11}$  is sent by Transmitter 1 along the  $3 \times 1$  beamforming vector  $(\mathbf{H}^{[21]})^{-1} \mathbf{S}^{[2]}$  so that when it goes through the channel  $\mathbf{H}^{[21]}$  to reach Receiver 2, where it is undesired, it will be rotated into the direction  $\mathbf{S}^{[2]}$ . The generic nature of channel matrices guarantees that any accidental alignments of the *desired* symbols do not take place. As a result, the desired symbols occupy linearly independent signal dimensions and can be isolated for a total of 4 DoF. Note that the corresponding two user interference channel setting allows only 3 DoF [65].

### 3.3 Interference Channel with $K > 2$ Users

The next setting for interference alignment was the  $K$  user interference channel [22]. This work further cemented the status of interference alignment as a general principle by establishing its applications in a variety of contexts ranging from MIMO beamforming solutions

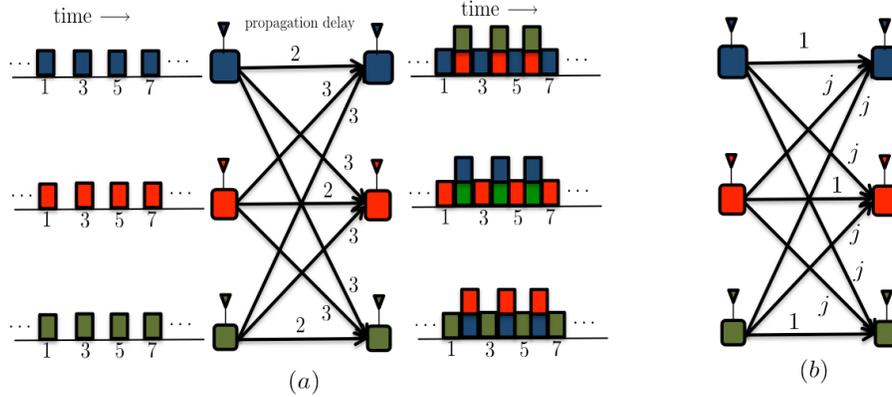


Fig. 3.3 Interference alignment on the 3 User Interference Channel [22] (a) through propagation delays. Each transmitter is active half the time — over odd time slots. Desired links introduce even delays while interfering links introduce odd delays so that interference is aligned over even time slots at each receiver and desired symbols are received free from interference. (b) through phase alignment. Each transmitter sends only real signals (sacrificing the imaginary baseband dimension). The direct channels produce no phase shift while the interfering channels rotate the phase by  $90^\circ$ . All interference aligns in the imaginary dimension at each receiver and desired signals are recovered free from interference over the real dimension.

(Figure 3.4) to propagation delay based alignment (Figure 3.3(a)) and phase alignments (Figure 3.3(b)).

While the propagation delay based alignment shown in Figure 3.3(a) appears to be a very specialized setting, it is quite insightful. It is shown in [64] that even with random node positions (i.e., random propagation delays) and arbitrary number of users it is possible (almost surely) to achieve propagation delay based interference alignment such that every user can access the channel free from interference for a fraction  $1/2 - \epsilon$  of the time, where  $\epsilon > 0$  can be chosen to be arbitrarily small, by choosing the basic symbol duration small enough. The need to choose the symbol duration small enough, i.e., the need to have essentially unlimited resolvability of received signals, or equivalently, the need to make bandwidth arbitrarily large, is an inescapable feature, in one form or another, of almost all known  $K$  user interference channel alignment schemes that guarantee *half the cake* per user. In other words, each user is guaranteed half the cake, only if the cake can be made large enough. This phenomenon, called *bandwidth scaling*, is explicitly characterized

by Gropop et al. in [55] and its necessity is highlighted, also in a propagation delay based framework. Specifically, Gropop et al. show that in a  $K$  user interference channel, bandwidth should scale as  $O(K^{2K^2})$  and point out that this bandwidth scaling factor is evident in most interference alignment schemes for the  $K$  user interference channel. Propagation delay based interference alignment for  $K$  user interference channel is also investigated in [145] where, again,  $K/2$  DoF are shown to be achievable. References [64, 96] also explore the geometrical placements of nodes that allows perfect propagation delay based alignment schemes under a line of sight propagation model. Similar to the propagation delay example of Figure 3.3(a), the phase alignment example shown in Figure 3.3(b) appears to be another toy example, but carries a very useful insight leading to the asymmetric complex signaling approach that will also be presented later in this monograph. It is remarkable that the propagation delay example exploits only the knowledge of the propagation delays and does not require any knowledge of the channel fading coefficient values at the transmitters, while the phase alignment example exploits only the phase of the channel coefficients and does not require any knowledge of the strengths of the channel coefficients at the transmitters in order to align interference. Another interesting observation is that essentially the two models are similar, because both channel phase differences and channel propagation delay differences arise out of differences in channel propagation path lengths.

Relative to all previous applications of interference alignment, the interference channel setting with three or more users presents a fundamental new challenge in that each signal needs to satisfy more than one alignment condition. For instance, in a three user interference channel, Transmitter 1's signal needs to align with Transmitter 2's signal at Receiver 3, and with Transmitter 3's signal at Receiver 2. This creates a chain of alignment conditions that in the end folds back upon itself, as shown in Figure 3.4, so that the optimal signaling basis vectors end up being the basis vectors of an invariant subspace (e.g., eigenvectors in the example illustrated in Figure 3.4) of a product form of channel matrices.

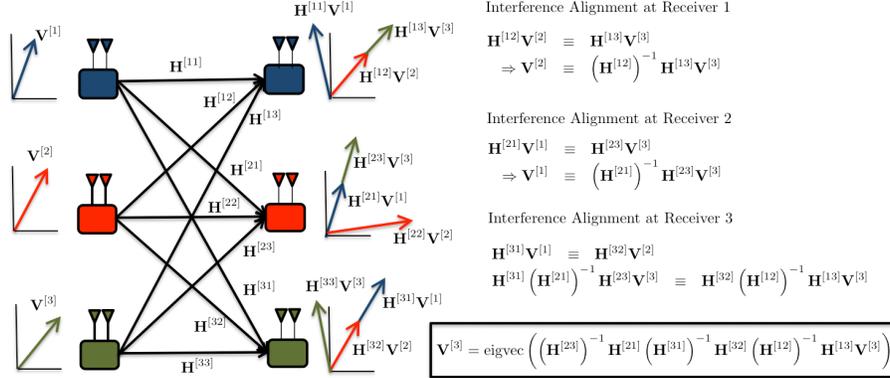


Fig. 3.4 Interference alignment on the 3 User MIMO Interference Channel with 2 antenna nodes [22] — Each receiver sees 2 linear equations in 3 unknowns. Since only one of the unknowns is desired, interference alignment makes it possible to recover its value by aligning the two undesired symbols along the same dimension at each receiver.  $A \equiv B$  is short hand notation for  $\text{column-span}(A) = \text{column-span}(B)$ .

For the example of Figure 3.4, interference alignment allows a total of 3 DoF (which is also the maximum value possible for this channel), whereas conventional schemes that orthogonalize the users can allow no more than a total of 2 DoF per channel use. Whereas in an orthogonalized solution, e.g., TDMA, each user would send two symbols on his  $2 \times 2$  MIMO channel and only one user will be active at a time, the interference alignment solution allows each user to send only one information symbol (thus sacrificing half the signaling dimensions available to that user) but all the users are active at all times.

# 4

---

## New Challenges and Solutions

---

The example of Figure 3.4 extends in a relatively straightforward manner to the case where each user is equipped with more than two antennas [22]. However, the example of Figure 3.4 does not extend to more than three users or even the three user setting when all nodes are equipped with only one antenna. In fact, the problem becomes *much* more difficult as the number of users increases beyond three or the number of antennas at each node is reduced to one. Two main issues faced by interference alignment schemes are:

1. *The number of alignment constraints* grows very rapidly as the number of interfering users is increased. For instance, in a  $K$  user interference channel, each of the  $K$  receivers needs an alignment of  $K - 1$  interfering signal spaces, for a total of  $O(K^2)$  signal space alignment constraints. Since there are only  $K$  signal spaces (one at each transmitter) to be chosen in order to satisfy  $O(K^2)$  signal space alignment constraints, the problem can quickly appear infeasible.
2. *The diversity of channels* which enables the relativity of alignment — which in turn is the enabling premise for

interference alignment — is often a limiting factor, e.g., when each node has only one antenna and all channels are constant across time and frequency. Limited diversity imposes fundamental limitations on the extent to which interference can be aligned in a network.

Further issues to be dealt with by interference alignment schemes include the imperfect, noisy, localized and possibly delayed nature of channel knowledge feedback to the transmitters where such knowledge is crucial to achieve interference alignment. These challenges have given rise to ingenious solutions in the form of new interference alignment schemes, invariably providing new and surprising insights into the non-trivial interplay of signaling dimensions in an interference network. The purpose of this section is to summarize these challenges and the corresponding solutions.

#### 4.1 Feasibility of Linear Interference Alignment

Linear interference alignment, i.e., alignment of signal spaces based on linear precoding (beamforming) schemes is the simplest form of interference alignment. Since beamforming schemes are common in existing MIMO point to point, BC and multiple access (MAC) networks, linear interference alignment is also the most easily accessible form of interference alignment from a practical perspective.

The feasibility of linear interference alignment is investigated for the MIMO interference channel setting in [48, 116, 124, 146, 162]. Following the notation of [162], let us define a  $(M_1 \times N_1, d_1)(M_2 \times N_2, d_2) \cdots (M_K \times N_K, d_K)$  interference alignment setting as an interference channel with  $K$  users where Transmitter  $k$ , equipped with  $M_k$  antennas, wants to send  $d_k$  independent streams of information symbols, each corresponding to one DoF, to Receiver  $k$  which is equipped with  $N_k$  antennas,  $k \in \mathcal{K} \triangleq \{1, 2, \dots, K\}$ . In the symmetric setting, i.e.,  $M_k = M, N_k = N, d_k = d$  for all  $k \in \mathcal{K}$ , this is denoted as an  $(M \times N, d)^K$  interference channel. Thus, Figure 3.4 presents the solution for the  $(2 \times 2, 1)^3$  interference channel.

The  $(2 \times 2, 1)^3$  setting is not the only setting where nontrivial *analytical* interference alignment solutions can be found. Reference [146] provides analytical achievable schemes for  $(N \times N, 1)^{N+1}$  systems. In some cases, with more constrained settings (i.e., with fewer free variables but the same number of alignment constraints), like  $(2 \times 3)^2(3 \times 2)^2$ , i.e., a four user interference channel with two  $2 \times 3$  users and two  $3 \times 2$  users, reference [162] finds analytical interference alignment solutions as well.

In general, the feasibility of a linear interference alignment solution for the  $(M_1 \times N_1, d_1)(M_2 \times N_2, d_2) \cdots (M_K \times N_K, d_K)$  interference channel is shown in [48, 162] to be equivalent to the existence of  $M_k \times d_k$  transmit precoding matrices  $\mathbf{V}_k$ ,  $\text{rank}(\mathbf{V}_k) = d_k$ , and  $N_k \times d_k$  receive combining matrices  $\mathbf{U}_k$ ,  $\text{rank}(\mathbf{U}_k) = d_k$ , such that for all  $j, k \in \mathcal{K}, j \neq k$ ,

$$\mathbf{U}_j^\dagger \mathbf{H}^{[jk]} \mathbf{V}_k = \mathbf{0}_{d_j \times d_k} \quad (4.1)$$

$$\text{rank}(\mathbf{U}_k^\dagger \mathbf{H}^{[kk]} \mathbf{V}_k) = d_k, \quad (4.2)$$

where  $\mathbf{H}^{[jk]}$  is the  $N_k \times M_j$  linear transformation representing the channel between Receiver  $j$  and Transmitter  $k$ . Of the two conditions, (4.1) may be interpreted as the condition for the existence of an interference free space of desired dimensions, while (4.2) is the condition that ensures that the desired signal is visible and resolvable within the interference-free space.

*Reciprocity:* An immediate consequence of this formulation is the reciprocity of linear interference alignment — linear interference alignment is feasible for the  $(M_1 \times N_1, d_1)(M_2 \times N_2, d_2) \cdots (M_K \times N_K, d_K)$  interference channel if and only if it is feasible for the  $(N_1 \times M_1, d_1)(N_2 \times M_2, d_2) \cdots (N_K \times M_K, d_K)$  interference channel, i.e., the reciprocal channel obtained by switching the roles of all transmitters and receivers. The reciprocity is evident by a corresponding switching of transmit precoding matrices  $\mathbf{V}_k$  and the receive combining matrices  $\mathbf{U}_k$ .

Conditions (4.2) and (4.1) are necessary and sufficient for feasibility of linear interference alignment regardless of whether the channel matrices are generic or arbitrarily chosen, and whether or not they

possess any specialized structure, such as (block) diagonal forms. Not much is known about the feasibility of linear interference alignment for arbitrary choice of channel coefficients. In fact, determining the feasibility of linear interference alignment for arbitrary channel realizations has been shown to be NP-hard by Razaviyayn et al. in [124] as the number of users becomes large. This is because for some channel realizations the feasibility problem subsumes the maximum independent set problem and the 3-colorability problem, both of which are known to be NP-complete. However, there has been substantial progress in the understanding of the feasibility of linear interference alignment over generic, structureless channels. Feasibility results for generic channels are invariably in the form of probabilistic guarantees that the alignment problem is feasible (infeasible) for *almost all* channel realizations, when the channels are drawn from a continuous distribution. Since wireless channels are determined by nature, and therefore generic, there is much interest in such feasibility characterizations. It is important to mention that generic channels can still have structure, e.g., block diagonal forms resulting from coding across parallel channel realizations (different time slots or carriers). While this structure has obvious practical relevance, the feasibility studies have so far been mostly limited to structureless generic channels. From a practical standpoint, structureless generic channels correspond to *spatial* precoding over multiple antennas *without* allowing symbol extensions, i.e., without coding across multiple channel realizations in time/frequency. The understanding of feasibility of linear interference alignment over generic structureless channels is reviewed next.

For the most commonly studied interference channel setting over single-hop wireless networks, the direct channels  $\mathbf{H}^{[kk]}$  are sufficiently *independent* of cross-channels  $\mathbf{H}^{[jk]}$  and have enough DoF ( $\text{rank}(\mathbf{H}^{[kk]}) = \min(M_k, N_k) \geq d_k$ ) so that the condition (4.2) is automatically satisfied *almost surely* if  $\mathbf{V}_k$  and  $\mathbf{U}_k$  with  $\text{rank}(\mathbf{V}_k) = \text{rank}(\mathbf{U}_k) = d_k$  are found only based on (4.1) with no regard to direct channels  $\mathbf{H}^{[kk]}$ . Under this assumption of independence of direct channels, the feasibility of linear interference alignment is equivalent to the existence of nontrivial solutions to the system of polynomial equations (4.1) alone.

The solvability of a system of polynomial equations is, in general, an open problem in algebraic geometry. However, it is known that a *generic* system of polynomial equations is solvable if and only if the number of equations does not exceed the number of variables. While the qualification of *generic* property is intended in a precise technical sense that does not *strictly* apply to the problem at hand, the intuition behind the statement is believed to be much more widely true. Relying on this intuition, reference [162] defines an interference alignment problem as improper or proper based on whether or not the number of equations exceeds the number of variables, with the understanding that proper systems are likely to be feasible and improper systems are likely to be infeasible. The upshot (Theorem 1 of [162]) is the qualification of the  $(M \times N, d)^K$  linear interference alignment problem as proper if and only if

$$d \leq \frac{M + N}{K + 1}. \quad (4.3)$$

It is notable that the  $(2 \times 2, 1)^3$  system of Figure 3.4 satisfies (4.3) with equality, i.e., all available variables are used to satisfy the alignment conditions. The reciprocity of linear alignment conditions is evident in the interchangeability of  $M, N$ . Another interesting observation here is that the number of antennas at each transmitter and receiver  $M, N$  affects the feasibility only through their sum  $M + N$ . Thus, for example, the systems  $(4 \times 1, 1)^4, (3 \times 2, 1)^2, (2 \times 3, 1)^4, (1 \times 4, 1)^4$ , which have the same value of  $M + N$  all satisfy the condition (4.3) with equality. While the feasibility of  $(4 \times 1, 1)^4$  and  $(1 \times 4, 1)^4$  is evident simply by zero-forcing at the transmitters and receivers, respectively, the feasibility of  $(3 \times 2, 1)^4, (2 \times 3, 1)^4$  is indicated by condition (4.3) and is verified numerically as well as analytically in [162].

In general, Yetis et al. label a system as proper if and only if every subset of alignment equations involves at least as many independent variables as the number of equations (which boils down to the condition (4.3) when every transmitter has the same number of antennas  $M$ , every receiver has the same number of antennas  $N$ , and every user wants the same number of Dof  $d$ ). While similar conclusions are also obtained by Tresch et al. in [146] and Negro et al. in [104],

rigorous connections between proper/improper systems and feasibility/infeasibility of linear interference alignment were initially only available for a few special cases. For example, Yetis et al. used Bernshtein’s theorem from algebraic geometry to analytically evaluate the feasibility of linear interference alignment because the number of solutions of the system of polynomials can be evaluated as the mixed volume of the Newton polytopes corresponding to the support (i.e., the monomials with non-zero coefficients) of the polynomials. Since the complexity of this calculation is prohibitive, its utility is limited to smaller dimensions, such as the  $(2 \times 3, 1)^4$  setting where reference [162] is able to analytically verify the feasibility of interference alignment. Again based on Bernshtein’s theorem, it was shown that improper systems are infeasible for the  $(M \times N, 1)^K$  setting. Interestingly, it was also shown that if  $d > 1$ , i.e., there are multiple beams per user, then proper systems satisfying condition (4.3) may still not be feasible. An example is the  $(3 \times 3, 2)^2$  setting which satisfies (4.3) but is not feasible [65].

The relationship conjectured by Yetis et al. in [162] between proper/improper systems and feasibility/infeasibility of linear interference alignment, is settled completely in one direction and partially in the other direction by Bresler et al. in [17] and by Razaviyayn et al. in [123]. Since the  $K = 2$  user case is completely solved in [65], these works focus on the  $K \geq 3$  setting. Specifically, it is shown that improper systems are indeed infeasible in general. In the other direction, it is shown that if all users desire the same DoF  $d$  and the number of antennas  $M_k, N_k$  are divisible by  $d$  for each  $k$ , then proper systems are also feasible. Further, in the symmetric  $(M \times N, d)^K$  setting, if both  $M, N$  are divisible by  $d$ , then the feasibility condition is precisely given by (4.3). If also  $M = N$  (square channel matrices), then the feasibility condition is precisely given by (4.3) even if  $M, N$  are not divisible by  $d$ . Note that since this claim is made specifically for  $K \geq 3$ , the  $(3 \times 3, 2)^2$  counter-example mentioned above is not applicable.

It is worthwhile to recall that these feasibility results are based on linear interference alignment with no symbol extensions. This is because symbol extensions create structure and dependencies in the channel matrices, which violates the assumption of generic channels. The significance of channel structure is highlighted by contrasting the

feasibility condition  $d \leq 2M/(K + 1)$  implied by (4.3) with the achievability of  $d = M/2$  DoF established in [22], both for the  $(M \times M, d)^K$  setting. While the former limits the total DoF to no more than 2 times the single user interference-free DoF, the latter implies unlimited total DoF as the number of users continues to increase.

A parallel body of work has focused on iterative algorithms to not only test the feasibility of linear interference alignment but also to numerically find the alignment solutions. This work is initiated by Gomadam et al. in [48] where the motivation for an iterative approach is to achieve interference alignment with only local channel knowledge, by exploiting the two-way nature of communication and the reciprocal nature of the physical propagation medium. The goal of achieving interference alignment leads to a fundamental departure from more conventional iterative signal space optimization approaches that seek interference avoidance. Interference avoidance is a selfish strategy where each user attempts to maximize his own signal-to-interference-and-noise power ratio (SINR) by trying to direct his desired signals away from the interference at his desired receiver. Especially at high SNR, the selfish approach leads to very suboptimal Nash equilibrium points. Instead, the interference alignment approach is primarily driven by goal of minimizing the size of the interference space at undesired receivers with little regard to the fate of the desired signal at the desired receiver. The new approach, combined with the reciprocity of the wireless propagation medium, leads naturally to the distributed and iterative interference alignment algorithm proposed in [48]. Starting with an arbitrary choice of transmit precoding vectors  $\mathbf{V}_k$ , the goal is for each receiver to identify the directions along which it receives the least interference. Then, as the transmitters and receivers switch roles in the reciprocal network, the new transmit precoding vectors are same directions along which least interference was received in the original network. The main insight is that the directions along which the least interference is seen at a receiver from its undesired transmitters in the original network will also be the transmit directions that cause the least interference at undesired receivers in the reciprocal network. This leads to an iterative algorithm that primarily seeks to consolidate interference into a small number of dimensions. Due to the nonconvex nature of the problem

convergence to the global optimum is not guaranteed. Since interference alignment solutions are optimal primarily at high SNR, Gomadam et al. [48] also propose another iterative algorithm, called the *max-SINR* algorithm, which combines both the selfish and alignment objectives, in a simple albeit ad-hoc manner, making it difficult to analyze the convergence properties. As expected, the max SINR algorithm shows significant benefits at low to intermediate SNR values and approaches interference alignment at high SNR.

Following [48], several iterative algorithmic approaches have been proposed for the linear interference alignment problem. Reference [116] presents an *alternating minimization* approach that does not explicitly assume channel reciprocity, but is algorithmically identical to the distributed interference alignment algorithm of [48]. Reference [127] proposes an alternative approach based on weighted minimum mean-square error (MMSE) beamforming which leads to a better-behaved objective, compares favorably to the max-SINR algorithm and can also accommodate unequal priorities for the users' rates. MMSE based iterative algorithms and better behaved versions of the max-SINR algorithm are also presented in [118, 131]. The optimality of the max-SINR algorithm within the class of linear beamforming algorithms and its local convergence with exponential rate, are established at high SNR in [114]. Analytical approximations for the high-SNR slope and the SNR-offset for linear interference alignment schemes in large systems limit (number of users  $K$ , and the number of antennas  $N$  approach infinity) are presented by Schmidt et al. in [128] for the single-beam MIMO setting  $(N \times N, 1)^K$  when  $K = 2N - 1$  — where the feasibility condition (4.3) for linear interference alignment is tight — and compared with  $K = N$  when simple zero-forcing can be used. Max-SINR algorithm is found to outperform alignment and zero-forcing schemes.

A common limitation of algorithmic linear interference alignment approaches is the need to specify a-priori the values  $d_k$ , i.e., the number of signaling dimensions to be allocated to each user. The algorithms are designed to perform well only if the desired allocation is feasible, failing which no performance optimization is attempted. Recent progress in this direction comes from [112] which seeks to maximize the number of accessible dimensions even when the requested allocation is infeasible.

Another interesting approach comes from reference [124] which circumvents the problem of choosing  $d_k$  by optimizing directly over the users' input covariance matrices instead, and is shown to perform well in terms of data rates.

In general, the optimization of linear precoders and combining filters to maximize the sum-rate in an interference network, remains very much an open problem of great practical significance, especially at intermediate values of SNR. The problem is made even more challenging in light of the potential benefits of symbol extensions, inseparability of parallel channels and asymmetric complex signaling schemes, all of which fall within the class of linear beamforming schemes. In the next few sections, we explore each of these ideas in detail.

## 4.2 Symbol Extensions

Linear interference alignment schemes, i.e., those based on spatial beamforming, operate within the spatial dimensions provided by multiple antennas at the transmitting and receiving nodes, and seek to divide those spatial dimensions into separable subspaces to be occupied by interference and desired signals at each receiver. When the number of antennas at each node is insufficient, e.g., when each node is equipped with only one antenna, spatial interference alignment schemes do not find a vector space large enough to function. Further, since the number of beams must be an integer, purely spatial beamforming based interference alignment schemes can only achieve an integral number of signal dimensions per message per channel use. This was initially not seen as a limitation, especially since the maximum number of signaling dimensions (DoF) for all previously considered wireless networks (MIMO point to point [41, 142], MAC [147], BC [151, 152, 165], and two user interference channels [65]) were known to be integer values. The point to point MIMO channel with  $M$  transmit,  $N$  receive antennas has  $\min(M, N)$  DoF, the two user MIMO MAC with  $M_1, M_2$  transmit antennas and  $N$  receive antennas has  $\min(M_1 + M_2, N)$  DoF, the two user MIMO BC with  $M$  transmit antennas and  $N_1, N_2$  receive antennas has  $\min(M, N_1 + N_2)$  DoF, and the two user MIMO Interference Channel with  $M_1, M_2$  antennas at

the transmitters and  $N_1, N_2$  antennas at their corresponding receivers, has  $\min(M_1 + M_2, N_1 + N_2, \max(M_1, N_2), \max(M_2, N_1))$  DoF, all of which are integer values and all of which can be achieved by spatial beamforming.

The idea of using symbol extensions, i.e., beamforming across multiple channel uses, was introduced in [67] for the two user MIMO  $X$  channel where all nodes are equipped with  $M > 1$  antennas. Using symbol extensions [67] enlarged the achievable DoF of the two user MIMO  $X$  channel from the  $\lfloor \frac{4}{3}M \rfloor$  DoF previously achieved by [92] to  $\frac{4}{3}M$ , which also matched the outer bound for this class of channels. Thus, for example, the two user  $X$  channel with  $M = 2$  antennas at each node was shown to have exactly  $8/3$  DoF, and achievability was established by using linear beamforming across three channel uses. Specifically, over three channel uses, the size of the total signal space at each node is increased to 6 dimensions. Two independent symbols are sent for each of the four independent messages in the two user  $X$  channel. At each receiver, the four desired symbols (two from each transmitter) occupy a total of four signal dimensions while the remaining two signal dimensions are set aside for aligned interference from both transmitters. Since a total of eight information symbols are sent over three channel uses in a manner that allows each symbol to be resolved from interference at its desired receiver, the achieved DoF equal  $8/3$ . In parallel with the  $X$  channel [67], the idea of beamforming over symbol extensions was also introduced in the MISO compound BC setting in [159].

### 4.3 Asymmetric Complex Signaling

Symbol extensions increase the dimensionality of the vector space making it possible to access fractional signal dimensions through beamforming schemes, e.g., a one dimensional signal sent over three channel uses occupies only one-third signal dimensions. However, symbol extensions over constant channels do not automatically provide the diversity of linear transformations that is needed for linear interference alignment. This is especially evident for networks with only single antenna nodes. Symbol extensions over constant channel values in this setting only create channel matrices that are scaled identity matrices. Thus,

any beamforming vector  $\mathbf{V}$  sent by the transmitter in this larger space maintains an identical direction at each receiver, differing from one receiver to another only by a scaling factor. The lack of rotations ensures that the alignment of vector spaces is identical at each receiver, thus preventing the relativity of alignment necessary for linear interference alignment schemes. When dealing with complex numbers, i.e., when all channel coefficients, transmitted and received symbols as well as the noise are complex variables, a remedy to find distinct rotations is to exploit the phase rotations, i.e., the distinct phase-shifts inherent in the multiplication of complex scalars which can be seen as rotations in the two-dimensional real-imaginary plane. This is the idea of *asymmetric complex signaling*, introduced in [27]. We summarize the idea below using the X channel as an example.

Consider the two user X channel where each node is equipped with only  $M = 1$  antenna. This channel has  $\frac{4}{3}$  DoF. The solution for the case where each node is equipped with  $M = 3$  antennas is illustrated in Figure 3.2. With only one antenna at each node the first thought, especially in light of the preceding discussion on symbol extensions, would be to use linear beamforming over three symbols. This gives us a three dimensional space at each node, much like the case where each node is equipped with three antennas. However, this does not immediately solve the problem. The resulting  $3 \times 3$  channel matrix  $\mathbf{H}^{[rt]} = h^{[rt]}\mathbf{I}$ , i.e., each channel matrix is a scaled version of the identity matrix. Following Figure 3.2 the received signals at Receiver 1 can be expressed as:

$$\begin{aligned} & \mathbf{H}^{[12]}(\mathbf{H}^{[22]})^{-1}\mathbf{S}^{[2]}x_{12} + \mathbf{H}^{[11]}(\mathbf{H}^{[21]})^{-1}\mathbf{S}^{[2]}x_{11} + \mathbf{S}^{[1]}(x_{22} + x_{21}) \\ &= \mathbf{S}^{[2]} \left( \frac{h^{[12]}}{h^{[22]}}x_{12} + \frac{h^{[11]}}{h^{[21]}}x_{11} \right) + \mathbf{S}^{[1]}(x_{22} + x_{21}), \end{aligned}$$

i.e., both desired signals are aligned along  $\mathbf{S}^{[2]}$  and both undesired signals are aligned along  $\mathbf{S}^{[1]}$ . Since the two desired signals occupy the same dimension, this approach fails. To see how this problem can be avoided, consider an information symbol, say  $x$ , that is being transmitted along a beamforming vector, say  $\mathbf{V}$ , so that the transmitted signal is  $\mathbf{V}x$ . Since all symbols are complex, and each complex dimension can

be seen as two real dimensions, let us think of  $x = x_1 + jx_2$ , where  $x_1$  and  $x_2$  are the two real components of complex symbol  $x$ , respectively. Suppose  $x_1$  and  $x_2$  are independently encoded symbols. Then both these symbols are being sent along the same beamforming vector  $\mathbf{V}$ , an artifact of the common assumption that transmitted signals are proper (circularly symmetric) Gaussians or, more generally, are drawn from symmetric constellations. Instead, let us choose different beamforming vectors for  $x_1$  and  $x_2$ , so that the transmitted signal is  $\mathbf{V}_1x_1 + \mathbf{V}_2x_2$ . This is the key to *asymmetric* complex signaling.

With this new insight, we interpret the three complex signaling dimensions obtained by symbol extensions as 6 *real* signaling dimensions and remove certain unnecessary symmetries. Specifically, pick two interference alignment vectors at each receiver, so that Receiver 1 requires interference to align along  $\mathbf{S}_1^{[1]}, \mathbf{S}_2^{[1]}$  and Receiver 2 requires interference to align along  $\mathbf{S}_1^{[2]}, \mathbf{S}_2^{[2]}$ . Each  $\mathbf{S}_j^{[i]}$  is a  $3 \times 1$  vector with complex elements or equivalently a  $6 \times 1$  vector with real elements if we separate the real and imaginary parts and stack them vertically. All  $\mathbf{S}_j^{[i]}$  are chosen as generic vectors, i.e., generated from a continuous distribution, so that they are linearly independent almost surely. The transmitted signals are then composed as:

$$\begin{aligned} \text{Transmitter 1: } & \frac{1}{h^{[21]}}(\mathbf{S}_1^{[2]}x_{11,1} + \mathbf{S}_2^{[2]}x_{11,2}) \\ & + \frac{1}{h^{[11]}}(\mathbf{S}_1^{[1]}x_{21,1} + \mathbf{S}_2^{[1]}x_{21,2}) \end{aligned} \quad (4.4)$$

$$\begin{aligned} \text{Transmitter 2: } & \frac{1}{h^{[12]}}(\mathbf{S}_1^{[1]}x_{22,1} + \mathbf{S}_2^{[1]}x_{22,2}) \\ & + \frac{1}{h^{[22]}}(\mathbf{S}_1^{[2]}x_{12,1} + \mathbf{S}_2^{[2]}x_{12,2}). \end{aligned} \quad (4.5)$$

So that the received signal at, say Receiver 1, becomes:

$$\begin{aligned} & \mathbf{S}_1^{[1]}(x_{21,1} + x_{22,1}) + \mathbf{S}_2^{[1]}(x_{21,2} + x_{22,2}) + \mathbf{S}_1^{[2]} \left( \frac{h^{[11]}}{h^{[21]}}x_{11,1} + \frac{h^{[12]}}{h^{[22]}}x_{12,1} \right) \\ & + \mathbf{S}_2^{[2]} \left( \frac{h^{[11]}}{h^{[21]}}x_{11,2} + \frac{h^{[12]}}{h^{[22]}}x_{12,2} \right). \end{aligned} \quad (4.6)$$

Thus, out of the six real signaling dimensions available at Receiver 1, interference aligns along the two real signaling dimensions  $\mathbf{S}_1^{[1]}, \mathbf{S}_2^{[1]}$ . The remaining four desired real symbols can be resolved from the remaining four real signal dimensions if and only if:

$$\phi^{[11]} - \phi^{[21]} \neq \phi^{[12]} - \phi^{[22]} + n\pi \quad \text{for some } n \in \mathbb{Z}, \quad (4.7)$$

where  $h^{[ij]} = |h^{[ij]}|e^{j\phi^{[ij]}}$ , i.e.,  $\phi^{[ij]}$  is the phase of the channel coefficient  $h^{[ij]}$ . Thus, using asymmetric complex signaling it is shown in [27] that if (4.7) is satisfied — as it will be almost surely for generic channels — then the X channel with only one antenna at each node achieves 4/3 complex DoF (which is also the outer bound) with linear interference alignment.

Besides the X channel where it achieves the DoF outer bound, asymmetric complex signaling is used for the three user interference channel with single antenna nodes and constant channel coefficient values [27], and for the compound MISO BC in [52]. In general, while it increases the scope of linear interference alignment schemes, in many cases it does not suffice to achieve the DoF outer bound that might otherwise be achievable through more sophisticated interference alignment schemes. The limitation of asymmetric complex signaling comes from the observation that splitting a complex number into its real elements only provides us two dimensions within which signal vectors may be rotated. An important consequence of this is that if  $\mathbf{V}$  is some generic vector of complex coefficients and  $h_1, h_2$  are generic complex scalars, then as shown in [27],

1. The two vectors  $\mathbf{V}$  and  $h_1\mathbf{V}$  are almost surely *linearly independent over reals* (if  $\alpha_1\mathbf{V} + \alpha_2h_1\mathbf{V} = 0, \alpha_1, \alpha_2 \in \mathbb{R}$ , then  $\alpha_1, \alpha_2$  must be zero).
2. The three vectors  $\mathbf{V}, h_1\mathbf{V}, h_2\mathbf{V}$  are almost surely *not* linearly independent over reals.

Thus, a phase rotation through multiplication with a complex scalar provides a linearly independent dimension over reals, but any further phase rotations through multiplications with complex scalars only create linearly dependent dimensions. For the interference channel this

limitation translates into the constraint that *any signal vector can align with interference at no more than 1 unintended receiver* [27]. But even in the three user interference channel each signal vector is seen by two unintended receivers. To elaborate on this limitation further, consider signal vectors  $\mathbf{V}_1, \mathbf{V}_2, \mathbf{V}_3$  originating at transmitters 1, 2, 3, respectively, in a three user interference channel with constant, complex channel coefficients. Ideally, we would like for  $\mathbf{V}_2$  to align with  $\mathbf{V}_3$  at Receiver 1, and with  $\mathbf{V}_1$  at Receiver 3. Since the channel is simply a complex constant, the received versions of the transmitted vectors are simply complex scaled versions of the transmitted vectors. So, if the desired alignments were to occur, then  $\mathbf{V}_1, \mathbf{V}_2, \mathbf{V}_3$  must be the same vector up to scaling by complex effective channel coefficients. As mentioned above, three complex scaled versions of the same vector cannot be linearly independent over reals. This would create a problem at Receiver 2, because as far as Receiver 2 is concerned, the desired signal vector  $\mathbf{V}_2$  will almost surely lie in the span of the interference vectors  $\mathbf{V}_1$  and  $\mathbf{V}_3$ , and hence will not be resolvable.

As a consequence of the limitation described above, the maximum achievable DoF for the three user interference channel with single antenna nodes, using linear interference alignment and asymmetric complex signaling is limited to only  $6/5$ . To see intuitively how this fraction  $6/5$  comes about, let us reason as follows. Consider a signal vector originating at Transmitter 1. At his desired receiver, this symbol must remain resolvable, so it must occupy one signal dimension. At one of the two remaining receivers, the symbol can align with an interfering symbol, i.e., it can share the same dimension with another symbol, so that it occupies only 0.5 signal dimension. At the last remaining receiver, however, this symbol cannot align with interference because of the limitation of asymmetric complex signaling discussed above. Therefore it will occupy one signal dimension at the last receiver. Summing up the signal dimensions occupied at all receivers by this signal vector, we find that this one signal vector casts a shadow of 2.5 signal dimensions at the receivers. A symmetric replication of this argument would conclude that each signal vector casts a shadow of no less than 2.5 signal dimensions that can be uniquely attributed to it. Thus, the DoF per user cannot be more than  $1/2.5$ . Equivalently, the total DoF of three users cannot be more

than  $3/2.5$ , i.e.,  $6/5$ . While this is still interesting, since  $6/5 > 1$ , i.e., it beats any orthogonal access scheme, it falls short of the DoF outer bound of  $3/2$  for the three user interference channel.

There is a broader insight revealed by the asymmetric complex signaling scheme, which has to do with the issue of proper versus improper Gaussian input distributions as well as with the need for symbol extensions.<sup>1</sup> The optimization of linear precoders under the assumption of Gaussian codebooks and treating interference as noise is the subject of a number of works under a variety of settings ranging from MIMO point to point, MAC, BC, and interference channels. Invariably the assumption is that the codebooks are proper (circularly symmetric) Gaussian codebooks. Since proper Gaussians are completely characterized by a covariance matrix (as usual we assume zero mean input distributions without loss of generality), the linear precoding optimization problem translates into the input covariance matrix optimization. The design of beamforming vectors is subsumed into the input covariance matrix optimization as well. Along with the more general class of improper Gaussians, symbol extensions are mostly ignored. These assumptions are guided by the conventional wisdom that channel extensions and improper Gaussians do not help, which in turn may be attributed to the well-known optimality of proper Gaussians without the need for symbol extensions for capacity achieving schemes in MIMO point to point, MAC and BC networks, i.e., all cases where capacity is known. However, the new insight from asymmetric complex signaling [27] shows that for settings where interference alignment is a possibility, improper Gaussians, along with symbol extensions, can outperform proper Gaussians, within the context of linear precoding schemes where all interference is treated as noise. Improper Gaussians need both covariance and pseudo-covariance matrix for a complete characterization. Thus, jointly considering more than one channel use, and by optimizing the pseudo-covariance matrix in addition to the covariance matrix, the resulting linear precoding schemes significantly

---

<sup>1</sup>A complex  $n \times 1$  random vector  $\mathbf{A} = \mathbf{A}_R + j\mathbf{A}_I$  is said to be proper if its pseudo-covariance matrix,  $E[(\mathbf{A} - E[\mathbf{A}])(\mathbf{A} - E[\mathbf{A}])^T] = 0$ . Note the use of only transpose instead of conjugate-transpose in the definition of the pseudo-covariance matrix. Proper Gaussian vectors are also called circularly symmetric.

outperform the single-channel use transmit covariance matrix optimization schemes. To the best of our knowledge such optimizations have not been attempted, and offer an interesting and wide open research avenue.

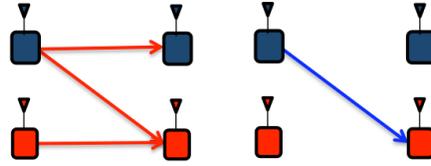
#### 4.4 Channel Variations: Separability

Channel variations — in both temporal and frequency dimensions — are an essential feature of wireless networks. The (ergodic) capacity of the time-varying point-to-point wireless channel is known to be the average of the capacities of the individual channel states subject to optimal power allocation (water-filling) across channel states [46]. The same is true for MIMO point to point, MAC and BC channels [45]. The general intuition from these results is that *the whole is equal to the sum of its parts*, i.e., the capacity of a varying channel can be achieved by achieving capacity individually for each fixed channel state. In other words, from a capacity perspective, there is no need to jointly code across channel states. This is known as the *separability* property of fading MIMO point to point, MAC and BC channels.

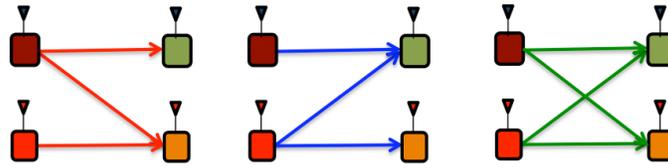
Starting with [24] several recent works [25, 34, 126] have established that the separability property is more of an exception than a rule for wireless networks in general and interference networks in particular. Two key causes for inseparability are highlighted in Figure 4.1(a) and Figure 4.1(b) and explained below. Note that we define a network to be sum-capacity-separable if and only if it is sum-capacity-separable for all possible realizations of channel states — i.e., if for all possible choices of channel coefficient values the sum capacity of the network can be expressed as the sum of the sum-capacities of channel states (subject to optimal power allocation across states).

##### *Causes for Inseparability*

1. *Antidote Links*: Figure 4.1(a) shows two inseparable states of the two user interference channel. In the blue channel state there are no links between desired transmitter–receiver pairs, which make the capacity of this state zero. However, the cross link can still be used to deliver the antidote to



(a) Two inseparable states of an interference channel



(b) Three inseparable states of an X channel

Fig. 4.1 Causes for inseparability — (a) antidote links (b) interference alignment. All channel coefficients shown are equal to unity for these two examples.

the interference that will be seen in the red channel state. Thus, the red and blue channel states that by themselves can only achieve 1 and 0 DoF, respectively, taken together can achieve 2 DoF, i.e., the whole is more than the sum of its parts. More generally, the example shows that the presence of links that cannot carry desired messages always makes a network inseparable. Remarkably the states of Figure 4.1(a) are *separable* from a sum-capacity perspective if the cross link can also carry a message, e.g., if we consider this to be an X channel instead of an interference channel [25].

2. *Interference Alignment*: Figure 4.1(b) shows three inseparable states of an X channel. In an X channel since there is a desired message from each transmitter to each receiver, all links can be used to send desired information. However, because of interference alignment this network is also not separable. While each state by itself has only 1 DoF (i.e., does not allow interference alignment), the three states together allow 4 DoF due to interference alignment [25]. Thus, once again the sum-capacity of the three states taken together is more than the sum of the sum-capacities of each.

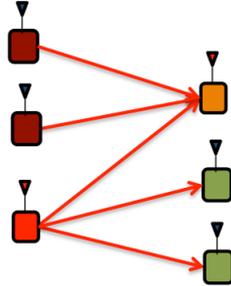


Fig. 4.2 The MAC-Z-BC network [25] is sum-capacity separable over any arbitrary set of states. It is the only one-hop wireless network to have this property.

Identifying these two causes of sum-capacity-inseparability allows reference [25] to show that all single hop wireless network topologies are sum-capacity-inseparable, *with the exception of the MAC-Z-BC network shown in Figure 4.2*. What makes the MAC-Z-BC network unique is that it is the only network that has no antidote links (there is a desired message from every transmitter to every receiver) and does not allow the possibility of interference alignment — since the BC channel from the red transmitter will serve only the stronger of the green receivers for sum-capacity, the maximum number of undesired messages for the orange receiver is only 1. The result is shown rigorously in [25] where also the sum-capacity of the MAC-Z-BC network is derived for arbitrary channel state values. Note that since the network is separable, the sum-capacity for the time-varying/frequency-selective setting is also automatically solved. The sum-capacity separability of the MAC, BC, and Z channels is also implied by the sum-capacity-separability of the MAC-Z-BC network because these networks can be simply produced by setting appropriate channel coefficient values to zero. It is also worth pointing out that even the MAC-Z-BC network is not separable from a capacity region perspective. From the entire capacity region perspective, the MAC and the BC are the only separable one-hop wireless networks.

The inseparability of wireless networks is significant in that finding the capacity of an inseparable interference network with fixed channel state would still leave the capacity of the same interference network under time/frequency selective conditions open. As it turns out this is the case for all wireless networks except the MAC-Z-BC network.

As shown by the X channel example in Figure 4.1(b), channel variations across different states play an important role in making interference alignment feasible even when it would be impossible over any given channel state. Incidentally, the first use of channel variations for interference alignment is also in the context of the X channel in [67]. In light of the previous section, the channel variations provide additional *diversity* which is very helpful in particular for linear interference alignment schemes. Even with single antenna nodes, symbol extensions across varying channel states create channel matrices that are diagonal but not scaled identity matrices. Since diagonal matrices do rotate the beamforming vectors, and these rotations are different for each transmitter–receiver pair, this creates the relativity of alignment needed for interference alignment.

An elegant study of the limited diversity provided by channel variations across a finite set of states, and its impact on the feasibility of linear interference alignment, is provided by Bresler and Tse in [20]. The key observation is that for generic channels  $\mathbf{H}_i$  of diversity  $D$  (e.g., signaling over multiple carriers with up to  $D/2$  independent carriers, each carrying two dimensional, i.e., complex signals) the vectors  $\mathbf{V}, \mathbf{H}_1\mathbf{V}, \mathbf{H}_2\mathbf{V}, \dots, \mathbf{H}_{N-1}\mathbf{V}$  are almost surely linearly independent over reals only for  $N \leq 2D - \lfloor D/2 \rfloor - 1$  and are almost surely dependent otherwise. For instance, the asymmetric complex setting discussed in the previous section can be seen as allowing diversity order 2, and therefore, as pointed out earlier,  $\mathbf{V}, h\mathbf{V}$  are almost surely linearly independent over reals while  $\mathbf{V}, h_1\mathbf{V}, h_2\mathbf{V}$  are almost surely linearly dependent over real dimensions. Bresler and Tse study in particular the three user interference channel with single antenna nodes where they show that with diversity order  $D$ , the maximum achievable DoF with linear interference alignment is limited to  $3D/2D + 1$ . As a sanity check, we note that this coincides with the 6/5 DoF limit observed with asymmetric complex signaling in [27] for the same setting (where  $D = 2$ ).

The inseparability of parallel interference networks is also quite significant from a practical perspective. For instance, OFDMA based multi-cell systems naturally correspond to parallel interference networks for each sub-carrier. Current resource allocation schemes follow a separation based approach where each subcarrier is treated separately.

However, due to the inseparability of parallel interference channels, performance improvements are possible through joint coding across sub-carriers. Da and Zhang show in [35] that these gains are most significant for cell-edge users who experience comparable signal strengths from their desired and interfering base stations, prompting a hybrid scheme whereby interference alignment is used for cell-edge users and conventional resource allocation is used for users closer to their desired base stations.

The topic of linear interference alignment over time-varying or frequency-selective channels is discussed in greater detail in the next two sections.

## 4.5 Ergodic Interference Alignment

Consider the following two states of a three user interference channel:

$$\mathbf{H}_a = \begin{bmatrix} 1 & -1 & 1 \\ 1 & 1 & -1 \\ -1 & 1 & 1 \end{bmatrix} \quad \mathbf{H}_b = \begin{bmatrix} 1 & 1 & -1 \\ -1 & 1 & 1 \\ 1 & -1 & 1 \end{bmatrix}, \quad (4.8)$$

where in each state, the  $(i, j)$ th element of the channel matrix represents the channel fading coefficient between Transmitter  $j$  and Receiver  $i$ . This is the original example presented in [24] to establish the inseparability of interference channels due to the possibility of interference alignment by coding across states even when each state by itself does not allow interference alignment. Specifically, if each transmitter is restricted to a power constraint  $P$  and the local noise power at each receiver is normalized to unity, then each channel state by itself has a sum capacity  $\log(1 + 3P)$  while the sum-capacity with joint coding across the two states is  $3\log(1 + 2P)$ . The key property here is the *complementary* nature of the two channel states, in that  $\mathbf{H}_a + \mathbf{H}_b = 2I$ . If each transmitter repeats the same symbol over the two channel states and each receiver adds the channel outputs from the two channel states, then all interference is automatically cancelled (and the desired signals combine coherently). Note that it is because all *interference is aligned* along the vector  $[1 \ -1]^T$  while the desired signal is received along the vector  $[1 \ 1]^T$  that the two can be separated.

This idea leads to an opportunistic interference alignment scheme proposed in [101] that pairs complementary channel states to naturally cancel interference by simple repetition coding across the complementary state pairs. Specifically, for each fading channel state  $\mathbf{H}$  there is a complementary channel state  $\mathbf{H}_c$  such that the desired channels stay the same while the undesired channels change polarity, i.e.,  $\mathbf{H}(i, i) = \mathbf{H}_c(i, i)$  and  $\mathbf{H}(j, k) = -\mathbf{H}_c(j, k), j \neq k$ . In [101], Nazer et al. show that for a broad class of channel distributions where the channel coefficients are drawn independently of each other and in time and satisfy a symmetry condition (that the channel probability density assigns equal weight to the two polarities associated with each channel fade value, i.e.,  $p_{H_{j,k}}(h) = p_{H_{j,k}}(-h)$ ), it is possible for each user to achieve the same rate as if he was allowed to transmit with no interference for one-half of the time, regardless of the number of users, and regardless of the SNR values (the latter is especially significant, since the benefits of many interference alignment schemes are only realized at asymptotically high SNR values). Since channel fades are drawn from a continuous distribution which makes the probability of any given channel state equal to zero, Nazer et al. first quantize the channel states and then use a strong typicality argument to show that if the channel fading process is ergodic then almost all quantized states are matched with complementary states, i.e., the fraction of unmatched states approaches zero. Because of the reliance on the ergodic fading property, this approach is called *ergodic interference alignment*.

The ergodic interference alignment scheme [101] shows that in a  $K$  user interference channel each user can simultaneously achieve a rate equal to his TDMA-1/2 rate, i.e., the rate he could achieve if allowed to transmit with no interference for half the time. It is shown in [69] that this is also the best rate possible in the presence of *equal strength* interferers, i.e., interference that is as strong as the desired signal. Specifically, [69] shows that equal strength interferers act as bottlenecks on the sum-capacity of the interference channel in the sense that the presence of a relatively small number of equal-strength interference carrying links is enough to limit the capacity per user to the TDMA-1/2 rate regardless of the strengths of the remaining interference-carrying links. Since ergodic interference alignment achieves this rate, it allows an

*exact* sum-capacity characterization for a number of  $K$  user interference network scenarios called *bottleneck states*. Further, [69] shows that in a dense network (with identical desired channel strengths) there are enough bottleneck links that the capacity per user converges in probability to the rate guaranteed by ergodic interference alignment. The result is strengthened to less symmetric forms of dense networks, e.g., with random node placements, in [1, 72].

Ergodic interference alignment schemes for settings with general message sets (e.g., the X channel) are presented in [102]. Ergodic interference alignment is the basis for the inner bounds on the  $n \times n$ -dimensional unicast and  $n \times 2^n$ -dimensional multicast capacity regions of dense wireless networks with arbitrary (deterministic) node placement, that are shown to be tight within  $O(\log(n))$  in [105]. Ergodic interference alignment schemes are also presented over finite field models in [70, 101].

While ergodic interference alignment schemes are remarkably simple from a coding perspective, and even optimal at any SNR for a broad class of settings, an important limitation of ergodic interference alignment is the delay involved in waiting for complementary states. This delay grows with the number of users as  $O(c^{K^2})$  and, to a certain extent, can be traded-off against the data rate by enlarging the set of acceptable complementary channel states, as shown in [73, 81].

## 4.6 An Asymptotic Interference Alignment Scheme — [CJ08]

The ergodic interference alignment scheme discussed in the previous section is an opportunistic scheme that exploits the existence of complementary channel states in equal proportions to achieve linear interference alignment. While this assumption applies to a broad variety of channel distributions, including widely studied Rayleigh fading models, it is not universally applicable. Arbitrary channel distributions, or even standard ones like Rician fading, do not satisfy the symmetric phase assumptions made by ergodic interference alignment. This leaves open the question — how can interference alignment be achieved with generic channel distributions, i.e., channel distributions with no special

guarantees except that they are well-behaved (continuous). This question brings us to an asymptotic interference alignment scheme proposed in [22], that is the topic of this section. Since there are many alignment schemes that have an asymptotic character (including the ergodic interference alignment scheme) in one dimension or another, we will identify the asymptotic interference alignment scheme of [22] as the [CJ08] scheme. This scheme is primarily of theoretical interest because of its strongly asymptotic character which limits its practical use. However, as we will see in the remainder of this monograph, it is one of the most powerful theoretical interference alignment constructions in the sense that

1. It allows us to satisfy an arbitrarily large number of alignment constraints without compromising on the fraction of signal space available to each user as the number of users increases (with very mild assumptions on the channel distributions).
2. It is applicable in both linear and nonlinear (real alignment over rational dimensions) forms.
3. It is applicable to a variety of scenarios ranging from K-user interference channels, X networks, compound BC channels, cellular networks, and network coding applications including the distributed data storage repair problem.

#### 4.6.1 Essential Construction of [CJ08] Asymptotic Alignment Scheme

The key to satisfying an arbitrarily large number of alignment constraints is a mechanism proposed in [22] by which a set is constructed to be almost invariant to an arbitrarily large number of linear transformations. We first explain this essential construction.

**Problem Statement:** *We are given an arbitrary number  $N$  of linear transformations  $\mathbf{T}_1, \mathbf{T}_2, \dots, \mathbf{T}_N$ . Our goal is to construct a set  $\mathbf{V}$  of finite cardinality. The linear transformations  $\mathbf{T}_i$  are defined to act elementwise on the set  $\mathbf{V} = \{v_1, v_2, \dots, v_{|\mathbf{V}|}\}$  to produce the set  $\mathbf{T}_i\mathbf{V} = \{\mathbf{T}_i v_1, \mathbf{T}_i v_2, \dots, \mathbf{T}_i v_{|\mathbf{V}|}\}$  and are assumed to be commutative,*

i.e.,  $\mathbf{T}_i \mathbf{T}_j v_k = \mathbf{T}_j \mathbf{T}_i v_k$ , but are otherwise generic.<sup>2</sup> Construct a set  $\mathbf{V}$  consisting of only nonzero elements, such that

$$\frac{|\mathcal{I}|}{|\mathbf{V}|} \rightarrow 1, \quad (4.9)$$

where

$$\mathcal{I} \triangleq \mathbf{V} \cup \mathbf{T}_1 \mathbf{V} \cup \mathbf{T}_2 \mathbf{V} \cup \dots \cup \mathbf{T}_N \mathbf{V} \quad (4.10)$$

and the “approaching 1” notation is used to mean that for any  $\epsilon > 0$  we should be able to find a set  $\mathbf{V}$  such that  $\frac{|\mathcal{I}|}{|\mathbf{V}|} < 1 + \epsilon$ .

### Solution — [CJ08] Scheme

Simply put, our goal is to construct  $\mathbf{V}$  so that

$$\mathbf{V} \approx \mathbf{T}_1 \mathbf{V} \approx \mathbf{T}_2 \mathbf{V} \approx \dots \approx \mathbf{T}_N \mathbf{V}, \quad (4.11)$$

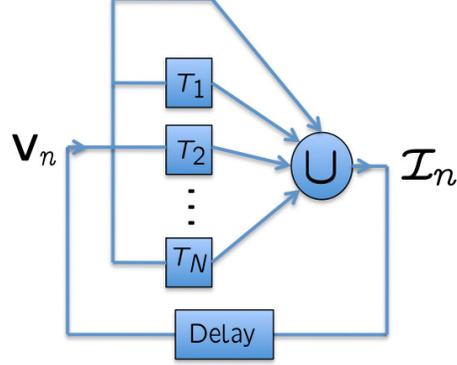
i.e., in an approximate sense,  $\mathbf{V}$  is invariant to the scaling factors  $\mathbf{T}_1, \dots, \mathbf{T}_N$ . In other words, all the linear transformations of the set  $\mathbf{V}$  align with the original set  $\mathbf{V}$ .

To gain intuition into the solution, let us start with a simple example where all linear transformations as well as the elements of  $\mathbf{V}$  are scalars. Suppose  $N = 1$ ,  $\mathbf{T}_1 = 2$ ,  $\mathbf{V} = \{1, 5, 7\}$ . Then we have  $\mathcal{I} = \{1, 5, 7, \dots\} \cup \{2, 10, 14\} = \{1, 2, 5, 7, 10, 14\}$ ,  $|\mathbf{V}| = 3$ ,  $|\mathcal{I}| = 6$ . Thus, an arbitrary choice of  $\mathbf{V}$  is, in general, not good.

Consider again the setting  $N = 1$ , an arbitrary  $\mathbf{T}_1$  and choose  $\mathbf{V} = \{w, \mathbf{T}_1 w, (\mathbf{T}_1)^2 w, \dots, (\mathbf{T}_1)^{n-1} w\}$ , where  $w$  is an arbitrary non zero scalar, so that  $\mathbf{T}_1 \mathbf{V} = \{\mathbf{T}_1 w, (\mathbf{T}_1)^2 w, \dots, (\mathbf{T}_1)^n w\}$  and  $\mathcal{I} = \{w, \mathbf{T}_1 w, (\mathbf{T}_1)^2 w, \dots, (\mathbf{T}_1)^n w\}$ . Now  $|\mathbf{V}| = n$ ,  $|\mathcal{I}| = n + 1$  and  $\frac{|\mathcal{I}|}{|\mathbf{V}|} = \frac{n+1}{n} \rightarrow 1$  by choosing  $n$  large enough. Evidently, the powers of  $\mathbf{T}_i$  are the key to the solution. With this intuition, we now consider the general problem.

The solution provided by the [CJ08] scheme is summarized in Figure 4.3 as an iterative loop with a unit delay. In a nutshell, starting with the set  $\mathbf{V}_1 = \{w\}$ , where  $w$  is a generic non zero element, we go

<sup>2</sup>By “otherwise generic,” here we mean that the transformations  $\mathbf{T}_i$  do not satisfy any other special property besides commutativity.

Fig. 4.3 The Construction of  $\mathbf{V}$  according to the [CJ08] scheme.

through the loop of Figure 4.3 a total of  $n$  times to obtain the set  $\mathbf{V}_n$  which satisfies (4.9) as  $n \rightarrow \infty$ . To see this explicitly, let us start with  $\mathbf{V}_1 = \mathbf{1}$ , i.e., a single column vector with all elements equal to unity. This produces the set

$$\mathcal{I}_1 = \{\mathbf{1}, \mathbf{T}_1\mathbf{1}, \mathbf{T}_2\mathbf{1}, \dots, \mathbf{T}_N\mathbf{1}\}. \quad (4.12)$$

Clearly at this point the alignment is poor. While the set  $\mathbf{V}_1$  has cardinality only 1, the set  $\mathcal{I}_1$  has  $N + 1$  elements.

Setting  $\mathbf{V}_2 = \mathcal{I}_1$ , we go through another iteration to obtain the new set

$$\mathcal{I}_2 = \{\mathbf{1}, \dots, \mathbf{T}_i\mathbf{1}, \dots, \mathbf{T}_i\mathbf{T}_j\mathbf{1}, \dots, \mathbf{T}_i^2\mathbf{1}\}. \quad (4.13)$$

In other words,  $\mathcal{I}_2$  contains all product terms of  $\mathbf{T}_i$  with overall degree  $\leq 2$ . The alignment is still poor. However, let us continue this iteration  $n$  times so that

$$\begin{aligned} \mathbf{V}_n = \{ & (\mathbf{T}_1)^{\alpha_1} (\mathbf{T}_2)^{\alpha_2} \dots (\mathbf{T}_N)^{\alpha_N} \mathbf{1}, \\ & \text{s. t. } \sum_{i=1}^N \alpha_i \leq n - 1, \\ & \alpha_1, \dots, \alpha_N \in \mathbb{Z}_+ \} \end{aligned} \quad (4.14)$$

and

$$\begin{aligned} \mathcal{I}_n = \{ & (\mathbf{T}_1)^{\alpha_1} (\mathbf{T}_2)^{\alpha_2} \dots (\mathbf{T}_N)^{\alpha_N} \mathbf{1}, \\ & \text{s. t. } \sum_{i=1}^N \alpha_i \leq n, \\ & \alpha_1, \dots, \alpha_N \in \mathbb{Z}_+ \} \end{aligned} \quad (4.15)$$

Thus,  $\mathbf{V}_n$  contains product terms up to degree  $n$ , while  $\mathcal{I}_n$  contains product terms up to degree  $n + 1$ . The number of column vectors in  $\mathbf{V}_n$  and  $\mathcal{I}_n$ , respectively, is

$$|\mathbf{V}_n| = \binom{n + N - 1}{N} \quad (4.16)$$

$$|\mathcal{I}_n| = \binom{n + N}{N}. \quad (4.17)$$

Therefore

$$\frac{|\mathbf{V}_n|}{|\mathcal{I}_n|} = \frac{n}{n + N} \rightarrow 1 \quad (4.18)$$

as  $n \rightarrow \infty$ . Since  $\mathbf{V}_n$  and  $\mathcal{I}_n$  have approximately the same size as  $n \rightarrow \infty$ , and since  $\mathbf{V}_n$  is always contained inside  $\mathcal{I}_n$ , we have achieved almost perfect interference alignment  $\mathcal{I} \approx \mathbf{V}$  asymptotically.

The key to the *alignment*,  $\mathcal{I} \approx \mathbf{V}$ , is the commutative property of the linear transformations  $\mathbf{T}_i$ . Without the commutative property,  $|\mathbf{V}_n|$  would not be able to catch up with  $|\mathcal{I}_n|$ . To see this, consider any term in  $\mathcal{I}_n$ . Starting with  $\mathbf{w}$ , this term can be seen as resulting from traversing the loop of Figure 4.3 a total of  $n$  times, each time choosing one of the  $N + 1$  parallel branches to cross over from  $\mathbf{V}$  to  $\mathcal{I}$ . The reason that the size of  $\mathcal{I}_n$  does not grow too quickly is because all permutations of the  $n$  choices made over this period produce the same output. For example, starting with  $\mathbf{w}$  if we choose  $\mathbf{T}_1$  the first time we traverse the loop,  $\mathbf{T}_2$  the second time, and  $\mathbf{T}_3$  the third time, the resulting term would be  $\mathbf{T}_3 \mathbf{T}_2 \mathbf{T}_1 \mathbf{w}$ . A different order of choices, say  $\mathbf{T}_2, \mathbf{T}_3, \mathbf{T}_1$  respectively, would not produce a distinct term only because the  $\mathbf{T}_i$  satisfy the commutative property. Thus, six different orders in which we could have chosen  $\mathbf{T}_1, \mathbf{T}_2, \mathbf{T}_3$  in traversing the loop three times, all

produce the same term in  $\mathcal{I}_3$ . Extending this argument to  $n$  iterations, it is clear that the commutative property is what keeps the size of  $\mathcal{I}_n$  from growing too quickly, so that ultimately,  $\mathbf{V}_n$  is able to catch up with it, and we achieve alignment  $\mathcal{I}_n \approx \mathbf{V}_n$  asymptotically as  $n \rightarrow \infty$ .

Next we explain how this essential construction is used to achieve interference alignment in the  $K$  user interference channel — the original setting of [22]. Other applications will be pointed out in detail in later sections.

#### 4.6.2 Application to $K$ -user Interference Channel with Time-varying Channel Coefficients

The original setting for the [CJ08] asymptotic interference alignment scheme is the time-varying or frequency-selective Gaussian interference channel where  $K$  transmitter–receiver pairs, labeled as users  $0, 1, \dots, K - 1$  wish to communicate  $K$  independent messages, one from each transmitter to its corresponding receiver, over the same wireless medium. For a detailed setup of this problem we refer the reader to [22]. For this review summary we highlight only certain interesting aspects and simplifications. With our notation, the received signal at Receiver 0 is expressed as (the description applies to both real and complex settings):

$$\mathbf{Y}^{[0]} = \mathbf{H}^{[00]}\mathbf{X}^{[0]} + \sum_{i=1}^{K-1} \mathbf{T}_i \mathbf{X}^{[i]} + \mathbf{Z}^{[0]}, \quad (4.19)$$

where  $\mathbf{H}^{[kk]}$  is the desired channel and  $\mathbf{T}_i$  enumerate the interference-carrying channels.  $\mathbf{Z}^{[k]} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  are the additive white Gaussian noise terms. Similarly, the received signal at Receiver 1 is expressed as:

$$\mathbf{Y}^{[1]} = \mathbf{H}^{[11]}\mathbf{X}^{[1]} + \sum_{i=1}^{K-1} \mathbf{T}_{K-1+i} \mathbf{X}^{[i+1]} + \mathbf{Z}^{[1]} \quad (4.20)$$

Note that all the superscripts in square parentheses are meant in the mod- $K$  sense. Thus, e.g.,  $\mathbf{X}^{[K]} = \mathbf{X}^{[0]}$ . Received signals at all other receivers are defined similarly.

The observant reader will note that there are a total of  $N = K(K - 1)$  normalized channels  $(\mathbf{T}_1, \mathbf{T}_2, \dots, \mathbf{T}_{K(K-1)})$  carrying

interference signals. Some simplification is possible here, as one could normalize one of the  $T_i$  for each receiver to the identity transformation and thereby reduce the number of variables. While such simplifications could improve the efficiency of the [CJ08] scheme, in the asymptotic regime of interest they will be inconsequential. Hence, we adopt the simpler convention of enumerating all the interference-carrying channels as linear transformations  $\mathbf{T}_i$ .

The interference alignment challenge here is to consolidate the interference carried by all these channels into a small subspace. In order to achieve the ultimate goal — that everyone gets half the cake<sup>3</sup> — we should be able to align all the interference at every receiver within one-half of the total signal space available at that receiver, leaving the other half interference-free for the desired signals. The number of alignment constraints is clearly increasing as the number of users and, therefore  $N$ , increases. Yet the DoF per user should remain unaffected — i.e., each user should continue to be able to access half the signal space free of interference. In other words, the number  $N$  is not important for the ultimate DoF result. One could potentially make  $N$  arbitrarily large, by adding more interference alignment constraints, e.g., by considering a finite state compound setting where each users' channel can take one out of finitely many generic values or by adding other fictional users without affecting the ultimate DoF-fraction available to each user. The alignment scheme will be able to satisfy an arbitrarily large number of alignment constraints while still guaranteeing that half the signal space seen by each receiver will be free from interference and available to the desired signal. This will be accomplished in an asymptotic sense by either enlarging the overall dimensionality of the signal space, or splitting signaling dimensions into sub-dimensions with increasingly fine resolution. In other words, the [CJ08] scheme guarantees that everyone will get half the cake almost surely, provided that either the cake can be made arbitrarily large, or that it can be cut with arbitrarily fine precision.

---

<sup>3</sup>The cake represents the capacity achieved by each user individually in the absence of all interference.

The most important assumption about the interference-carrying channels is the following.

- **The commutative property**, i.e.,  $\mathbf{T}_i\mathbf{T}_j = \mathbf{T}_j\mathbf{T}_i$ ,  $\forall i, j \in \{1, 2, \dots, N\}$ .

Note that this assumption is obviously satisfied by the diagonal channel matrices resulting from symbol extensions over time-varying/frequency-selective channels.

We will consider a beamforming scheme where each transmitter sends its information along  $m$  linearly independent beamforming vectors. These vectors can be viewed as the columns of an  $m \times 1$  precoding matrix  $\mathbf{V}$ . With some abuse of notation, we will use the symbol  $\mathbf{V}$  as a matrix, as a set (with the column vectors as its elements) and also as the signal subspace spanned by its columns. As a simplification of [22] we will make the following assumption.

- **All transmitters use the same set of signaling vectors  $\mathbf{V}$ .**
- **Each receiver will also set aside (approximately) the same subspace  $\mathbf{V}$  for interference.**

*Remark:* The approximation implied in this statement will be inconsequential in the asymptotic regime of interest.

With these assumptions in mind, let us account for the interference space at Receiver 0 (see Figure 4.4). The space associated with interference at Receiver 0 is the union of the space intentionally set aside for interference  $\mathbf{V}$  and the actual space along which interference shows up at Receiver 0.

$$\mathcal{I}^{[0]} = \mathbf{V} \cup \mathbf{T}_1\mathbf{V} \cup \mathbf{T}_2\mathbf{V} \cup \dots \cup \mathbf{T}_{K-1}\mathbf{V}. \quad (4.21)$$

This space is shown in Figure 4.5 as  $\mathcal{I}^{[0]}$ .

Let us collect all the interference space associated with *all the receivers* and express it as a larger subspace

$$\mathcal{I} = \mathbf{V} \cup \mathbf{T}_1\mathbf{V} \cup \mathbf{T}_2\mathbf{V} \cup \dots \cup \mathbf{T}_N\mathbf{V} \quad (4.22)$$

Then our interference alignment goal is expressed as:

$$\mathbf{V} \approx \mathbf{T}_1\mathbf{V} \approx \mathbf{T}_2\mathbf{V} \approx \dots \approx \mathbf{T}_N\mathbf{V} \quad (4.23)$$

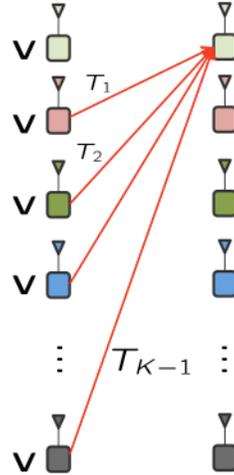


Fig. 4.4 Interference seen by Receiver 0.

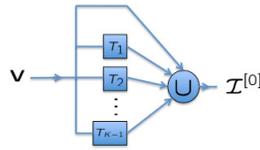


Fig. 4.5 Interference space at Receiver 0.

The [CJ08] IA scheme essentially aligns all the interference within the same space  $\mathbf{V}$  at each receiver. In other words, **the goal is to make  $\mathcal{I} \approx \mathbf{V}$** . From the discussion at the beginning of this section, we already know how to solve this problem.

While all interference is aligned, an important concern remains — that the desired signal must not overlap with the interference space. It may be especially concerning at first to note that each transmitter is precoding signals into the same signal space  $\mathbf{V}$  and  $\mathbf{V}$  is also the space along which interference aligns at each receiver. However, what keeps the desired signal distinct from interference is the *generic* linear transformation  $\mathbf{H}^{[k;k]}$  corresponding to the desired channel. Note that  $\mathbf{H}^{[k;k]}$  does not appear anywhere in the alignment problem described above, and is therefore *independent* of the alignment solution  $\mathbf{V}$ . Because of this independence the linear transformation  $\mathbf{H}^{[k;k]}$  essentially

transforms the desired signal for each receiver into *general position*  $\mathbf{H}^{[kk]}\mathbf{V}$  relative to the interference space  $\mathbf{V}$  seen by the receiver. It is well known that vector subspaces in general position do not overlap if they exist in an overall space that is large enough to accommodate all of them. Thus, all that is needed to ensure that the signal and interference are separable, is that the overall signal space be large enough to accommodate  $\mathbf{V}$  and  $\mathcal{I}$ . This is easily accomplished by choosing the overall signal space size to be  $S = |\mathbf{V}| + |\mathcal{I}|$ . With this setting,

$$\text{span}(\mathbf{H}^{[kk]}\mathbf{V}) \cap \text{span}(\mathbf{V}) = \{0\} \quad (4.24)$$

$$\frac{|\mathbf{H}^{[kk]}\mathbf{V}|}{S} = \frac{|\mathbf{V}|}{|\mathbf{V}| + |\mathcal{I}|} \rightarrow \frac{1}{2} \quad (4.25)$$

as  $n \rightarrow \infty$ , and thus everyone gets half the cake.

We summarize the interesting features of the alignment scheme again as follows.

1. The interference carrying channels  $\mathbf{T}_i$  are required to satisfy the commutative property. Beyond this, they could satisfy any special property and do not need to be generic.
2. The desired-signal carrying channels  $\mathbf{H}^{[kk]}$  only need to be sufficiently generic to transform the desired signals out of the interference space. In particular these channels are not required to satisfy the commutative property.
3. All transmitters use the same signal space  $\mathbf{V}$ .
4. All interference at each receiver is asymptotically aligned within (approximately) the same space  $\mathbf{V}$ .
5. The number of alignment constraints  $N$  can be arbitrarily large, and still almost perfect interference alignment is achieved asymptotically over a sufficiently expanded signal space.

Finally, to conclude this section, let us roughly estimate the signal space extension needed by the [CJ08] scheme to approach the limiting value of “half the cake” per user. Suppose each user wants to access a fraction of the cake that is within a small constant  $\delta$  of the ideal value  $1/2$ , i.e.,

we want

$$\frac{1}{2} - \delta \leq \frac{|\mathbf{V}_n|}{|\mathbf{V}_n| + |\mathcal{I}_n|} = \frac{1}{1 + \frac{|\mathcal{I}_n|}{|\mathbf{V}_n|}} = \frac{1}{1 + \frac{n+N}{n}}. \quad (4.26)$$

The smallest value of  $n$  that satisfies this condition is

$$n^* = \left\lceil N \left( \frac{1}{4\delta} - \frac{1}{2} \right) \right\rceil. \quad (4.27)$$

Thus, the total size of the signal space needed is  $|\mathcal{I}_{n^*}| + |\mathbf{V}_{n^*}| = \binom{n^*+N}{N} \left(1 + \frac{n^*}{n^*+N}\right)$ . Using Stirling's formula,

$$\begin{aligned} \left(\frac{n^*+N}{N}\right)^N &\leq \binom{n^*+N}{N} \leq \left(\frac{n^*+N}{N}\right)^N e^N \\ &\Rightarrow \left(\frac{1}{4\delta} + \frac{1}{2}\right)^{K(K-1)} \left(\frac{2}{1+2\delta}\right) \leq |\mathcal{I}_{n^*}| + |\mathbf{V}_{n^*}| \\ &\leq \left(\frac{e}{4\delta} + \frac{e}{2} + \frac{e}{2K(K-1)}\right)^{K(K-1)} \\ &\quad \times \left(\frac{2 + 4\delta/(K(K-1))}{1 + 2\delta + 2\delta/(K(K-1))}\right), \end{aligned} \quad (4.28)$$

where we made the substitution  $N = K(K-1)$ . Thus, for a fixed tolerance  $\delta$  around the target value of  $1/2$  DoF per user, as the number of users,  $K$ , grows, the *logarithm* of the size of the signal space needed by the [CJ08] scheme expands as  $\Theta(K^2)$ . This need for rapid bandwidth expansion is a sobering observation that underscores the strongly asymptotic character of the [CJ08] scheme. Everyone can get arbitrarily close to half of the cake, provided the size of the cake grows very rapidly with the number of users.

#### 4.7 Interference Alignment Based on Separability of Rationally Independent Dimensions

The [CJ08] asymptotic interference alignment scheme is originally introduced for the  $K$  user interference channel with time-varying or frequency-selective channel coefficients. The channel variations naturally give rise to diagonal channel matrices which satisfy the commutative property. For constant channels, especially with *real* channel

coefficients, the beamforming approach which separates signal spaces based on linear independence is not sufficient and the DoF problem initially remained open. Interference alignment over constant channels is shown to be possible through lattice alignment schemes in the work of Bresler et al. [18]. The idea is advanced further in [26, 39, 68, 136] leading most recently to a breakthrough by Motahari et al. in [99] where the framework for interference alignment along *rationally independent* dimensions is established. Because of its close connection to the vector space alignment framework discussed in the preceding sections, we first explore the rational dimensions framework in this section, and then return to the more distinct forms of lattice alignment in the next section.

In its details, the concept of rational dimensions introduced in [99], is far from a trivial extension of the well understood concept of linear independence, and is built upon recent generalizations of the Khintchine-Groshev theorem [12] from the field of Diophantine approximation theory. Yet, the idea of rationally independent dimensions is intuitively quite similar to the idea of linear independence of vectors. A set of vectors is considered linearly independent if no nontrivial linear combination of the vectors can produce the zero vector. Similarly, a set of scalars is considered to be rationally independent if no nontrivial rational combination of the scalars — i.e., a linear combination with only rational coefficients at least one of which must be nonzero — can add up to zero. Just as signals received along linearly independent vectors are separable, Motahari et al. showed that superimposed lattices scaled by rationally independent scalars were also separable.

The rational dimension framework of [99] also provides a mechanism to translate the [CJ08] asymptotic interference alignment solutions designed for time-varying channels directly to constant channel setting. Instead of linear precoding through beamforming vectors, lattice codes scaled by a set of rationally independent scalars are used for signaling. Because scalar channel coefficients for cross-channels (interference) naturally satisfy the commutative property, and the direct channel coefficients are assumed generic, the requirements of [CJ08] scheme are satisfied. Thus, within the new framework of rational dimensions, the [CJ08] scheme is again found to achieve almost perfect asymptotic

alignment, leading to the DoF characterizations for the  $K$  user interference channel for almost all values of channel coefficients.

Next we summarize the essential elements of the rational dimensions framework.

#### 4.7.1 The Rational Dimensions Framework

With the rational dimensions framework, transmitted signals are constructed as:

$$x = \lambda(v_1x_1 + v_2x_2 + \cdots + v_dx_d) \quad (4.29)$$

where  $v_i$  are the scalar precoding coefficients (whose values are chosen depending on channel coefficient values),  $|v_i|^2 = O(1)$ ,  $x_i \in \{0, 1, \dots, Q - 1\}$  are coded scalar  $Q$ -ary symbols, and  $\lambda$  is a scaling coefficient required to satisfy the power constraint. The precoding coefficients  $v_i$  are chosen to be rationally independent at the desired receivers so that the information streams can be separated along distinct rational dimensions. Rational independence of  $v_i$  is equivalent to the condition that

$$v_1x_1 + v_2x_2 + \cdots + v_dx_d = 0 \quad (4.30)$$

has only one solution for rational values of  $x_i$ , and that is the all-zero solution.

The separability of rational dimensions is a consequence of the (extensions of) Khintchine–Groshev theorem [10, 12] in Diophantine approximation theory. The theorem shows that the minimum distance of the constellation  $\lambda \sum_{i=1}^m v_i x_i$  when  $x_i \in \{0, 1, \dots, Q - 1\}$ , is at least  $\frac{\lambda}{O(Q^{m-1+\epsilon})}$  for almost all rationally independent  $v_1, v_2, \dots, v_m$  and any  $\epsilon > 0$  (extensions of the theorem show that the statement holds even when the  $v_i$  lie on a non-degenerate manifold as in the [CJ08] scheme.).

Choosing  $Q = O(P^{\frac{1-\epsilon}{2(m+\epsilon)}})$  and  $\lambda = \frac{P^{1/2}}{Q}$  ensures that the total transmit power is  $O(P)$ . With this setting, the minimum distance of the constellation is greater than the standard deviation of the noise, which guarantees the separability of coded information streams. The rate pre-log achieved per stream is

$$\lim_{P \rightarrow \infty} \frac{\log(Q)}{0.5 \log(P)} = \frac{1 - \epsilon}{m + \epsilon}. \quad (4.31)$$

Thus, the real scalar dimension, conventionally thought of as one DoF (or half a DoF if normalized according to the complex setting) is split into  $m$  rational dimensions, each carrying  $\approx 1/m$  DoF.

The rational dimensions framework enabled by Diophantine approximation theory, and in particular the Khintchine–Groshev theorem [12] is a powerful theoretical tool to find DoF characterizations for networks with constant channel coefficients. Originally proposed for only real channel coefficients, the framework has recently been extended to complex channel coefficients as well by Maddah-Ali in [93]. The relative ease with which interference alignment schemes based on conventional vector space dimensions can be translated into the rational dimensions framework is quite remarkable. The scheme also inherits some of the limitations of the diophantine approximation theory literature. For instance, while it is shown by Motahari et al. [99] that almost all channel coefficients are *good*, i.e., they allow the separability of data streams sent along fractional rational dimensions, it is not yet possible to identify a given channel state as *good* or bad. Thus, e.g., while we know that almost all states of the  $K$  user interference with constant channel coefficients achieve  $K/2$  DoF, for almost all channel states the DoF of that particular channel state cannot be determined from the current understanding of rational dimensions. This is mainly because the Khintchine–Groshev theorem that is central to the rational dimensions framework only claims to be applicable to almost all rationally independent signaling dimensions but provides no means of identifying whether it is applicable to a given choice of rationally independent signaling dimensions. Lastly, because the theorem only provides  $O(Q^2)$  guarantees in the minimum distance of the constellation, without identifying the constant scaling factor, it is not yet possible to gauge the performance at any finite SNR.

## 4.8 Lattice Alignment

The debate between random coding and structured coding approaches is as old as information theory itself. Remarkably, both theory and practice have switched sides on the issue of random versus structured codes. When Shannon theory advocated random coding, practical schemes

were exclusively focused on structured codes that could be decoded with a reasonable complexity. That debate was essentially won by theory as practical schemes like turbo codes were found to mimic random coding schemes and thereby approach theoretical limits with reasonable complexity. In the new debate, the situation is reversed. Recent theoretical advances advocate highly sophisticated structured codes while the achievable rates considered practical draw largely on basic random coding arguments. Perhaps the strongest case for structured codes, also the topic of this section, is the idea of lattice alignment.

The main motivation for lattice alignment is the observation that since the sum of lattice points (codewords) is also a lattice point (a valid codeword), it may be possible to decode the sum of codewords when they belong to the same lattice even if the individual codewords by themselves are not decodable. This is useful when the codewords only carry undesired information from interfering transmitters, so that the receiver does not need or wish to decode each codeword individually, but would like to decode the total effective interference (the sum of interfering codewords) to obtain a clean (interference-free) observed signal for its own desired codeword. This is particularly useful in a strong interference setting where decoding the interference would be beneficial before decoding the desired codeword. As such lattice alignment schemes are critical when the information is viewed as layered across signal levels such as may be seen through the linear deterministic models of Avestimehr et al. [3].

Lattice alignment, therefore, refers to the use of lattice codes in an interference network, with the lattices scaled in such a manner that the undesired signals at an interfered receiver arrive on the same lattice (so that the sum of interferers is decodable while the individual codewords are not), and at the same time the desired signal stands apart, i.e., does not occupy the same lattice, e.g., due to an inconsistent scaling factor.

The idea of lattice alignment is introduced by Bresler, Parekh and Tse in [18] for the many-to-one interference channel where only one receiver experiences interference, and is shown to achieve the capacity of this channel within a constant number of bits/sec/Hz (regardless of SNR and channel parameters). The layering of information in

signal levels is an essential aspect of this problem because of the different strengths of signals arriving from each interferer to the interfered receiver.

The first use of lattice alignment for fully connected interference channels appears in [26] as a series of examples of real channels with constant channel coefficients that can approach arbitrarily close to the  $K/2$  DoF outer bound. Etkin and Ordentlich [39] explore the DoF of fully connected  $K$  user interference channel with real and constant channel coefficients and discover surprising and fundamental new insights regarding the distinct behavior of channels with rational versus irrational channel coefficients. In particular, they find that (a) a fully connected interference channel with only rational channel coefficients has DoF bounded strictly away from  $K/2$ , and (b) that if all desired channel coefficients are irrational and all interference-carrying channel coefficients are rational, then the outer bound of  $K/2$  is achieved. Together, these results showed the discontinuity of DoF at all fully connected channels with rational coefficients. The first result is based on a novel outer bounding technique based on additive combinatorics and is very surprising as perhaps the first instance of an outer bound that is sensitive to rationality of channel coefficients. The second result of Etkin and Ordentlich highlights that if cross-links are all rational, then interference can be made to align on the same integer lattice, while the irrational nature of direct links ensures that the desired signals stand apart from the interference. The role of rationality in determining signal dimensions was developed further Motahari et al. in [99], leading to the rational dimensions framework described in the previous section. Another intriguing development in this direction is the casting of the DoF problem in terms of Renyi information dimensions by Wu et al. in [161]. Specifically, Wu et al. provide a general formula for the DoF of the  $K$ -user real-additive noise interference channel in terms of a single letter optimization of a linear combination of Renyi information dimensions, and show that this approach can recover and improve upon known results with simpler proofs.

While lattice alignment plays a role in determining the DoF in [26, 39, 99], perhaps the more unique aspects of lattice alignment become evident in going beyond DoF, to *generalized* DoF (GDoF, see

Appendix B) characterizations, or even further to capacity characterizations within a constant gap. It is in these settings that the signal level dimension plays a central role as some signals are stronger than others, naturally giving rise to a layered code structure (inspired by the deterministic approach of [3]) where some interference layers may be stronger (i.e., received at higher signal levels) than some desired signal layers. A constructive approach then finds it useful to *remove* interference from stronger layers by decoding and subtracting the aggregate interference before decoding weaker desired signals.

It is this ability to *decode* aggregate interference that sets lattice alignment apart from vector space interference alignment techniques. Note that interference alignment techniques based on signal vector spaces may suffice for DoF characterizations, especially if the channels are time-varying and sufficiently generic, but invariably these techniques do not suffice to obtain GDoF characterizations. This is because vector space alignment techniques typically sacrifice *entirely* the vector subspace within which interference is consolidated. Evidently when all signals are equally strong this sacrifice tends to be DoF optimal. However, if the aligned interference does not occupy all signal levels, as is the case if, e.g., all interferers are very strong, then it is inefficient from a GDoF perspective to sacrifice a signal vector dimension where the unoccupied signal *levels* could have been used to send the desired signal. This observation is reinforced in several works. Sridharan et al. [136] find a tighter characterization (than possible with random codes) of the very strong interference condition for a symmetric  $K$  user interference channel where all cross channels are identical, based on lattice alignment.

In [68], Jafar and Vishwanath find the GDoF characterization for a symmetric  $K$  user interference channel where all cross channels are equally strong. Interestingly, the GDoF per user in this setting are the same as the two user interference channel [19, 40], due to interference alignment based on lattice codes. Intuitively — once interference is aligned, facing one interferer is the same as facing an arbitrary number of interferers, especially for the symmetric model where all interferers are equally strong and each receiver sees the same picture with only the identities of desired and interference signals changed appropriately.

Bandemer et al. [7] find the  $K$  user interference channel GDoF with a cyclic interference model where adjacent interferers may have different strengths but the same pattern is repeated in a cyclical fashion. Saha and Berry in [125] study the sum rate of a class of three user deterministic interference channels, which they label as the “shoe-string” model, and where there is no interference between two of the users. Note that the solution of the deterministic channel model capacity problem is typically the precursor of the GDoF characterization of the corresponding Gaussian channel model, as evident from, e.g., [7, 19, 58, 68].

Huang et al. [58] find the GDoF characterization for a symmetric version of the  $2 \times 2$  X channel illustrated in Figure 4.6(a). As usual, the X channel contains four independent messages,  $W_{ij}$ ,  $i, j \in \{1, 2\}$ , with  $W_{ij}$  representing the message from Transmitter  $j$  to receiver  $i$ . Eliminating either the messages  $W_{12}, W_{21}$  or  $W_{11}, W_{22}$  reduces the X channel into one of its underlying interference channels with the two remaining messages. It is well known that the X channel has  $4/3$  DoF while the interference channel has only 1 DoF. The DoF setting implicitly assumes all channels are equally strong. The GDoF setting allows us to gauge the benefits of the X channel over the underlying interference channels when some channels are stronger/weaker than others. In other words, the problem is of interest primarily as a means of exploring the benefit of interference alignment enabled by the X channel setting relative to the underlying interference channels, whose GDoF are well recognized in the form of the familiar “W” curve of Etkin et al. [40] for the symmetric setting. There are two interference

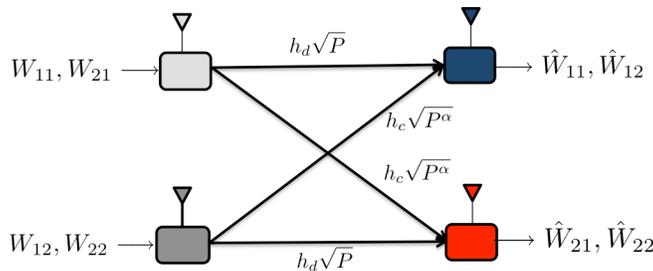


Fig. 4.6 X Channel extension of the Symmetric Interference Channel.  $W_{ij}$  is the message originating at Transmitter  $j$  and intended for Receiver  $i$ .

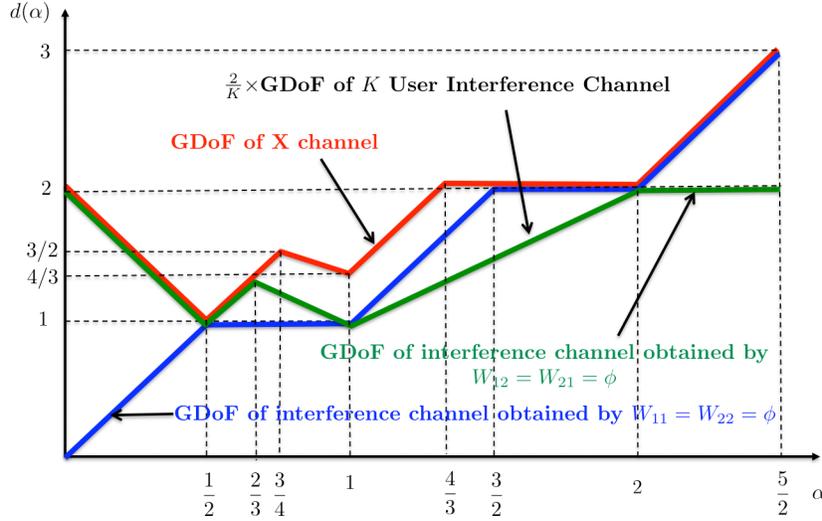


Fig. 4.7 GDoF of Symmetric Interference and X Channels [19, 40, 58, 68].

channels contained in the  $X$  channel. From Figure 4.6, it is clear that eliminating messages  $W_{12}, W_{21}$ , i.e., setting  $W_{12} = W_{21} = \phi$ , produces the interference channel with messages  $W_{11}, W_{22}$  where SNR values scale as  $P$  and INR values scale as  $P^\alpha$ . The GDoF curve is given by the recognizable "W" curve of Etkin et al. as shown in Figure 4.7. Alternatively, eliminating messages  $W_{11}, W_{22}$ , i.e., setting  $W_{11} = W_{22} = \phi$ , produces the interference channel with messages  $W_{12}, W_{21}$  where SNR scales as  $P^\alpha$  and INR scales as  $P$ . The GDoF curve for this interference channel is also obtained from the "W" curve of Etkin et al. with the suitable normalization of SNR and INR scaling factors, and is shown in Figure 4.7. As evident from Figure 4.7, the symmetric  $X$  channel retains a GDoF advantage over the corresponding symmetric two user interference channels in the range  $2/3 \leq \alpha \leq 3/2$  because of interference alignment. Another interesting observation from Figure 4.7 is that out of the two interference channels contained within the  $X$  channel, the channel under weak interference always outperforms the channel under strong interference in terms of GDoF. Intuitively, it means that instead of using the stronger channel to send interference that can be decoded and cancelled, the better use of the stronger channel is to send data instead, even if it means that the weaker channel will carry

interference that cannot be removed. As mentioned earlier, for all these results the key to the GDoF characterization is to solve the deterministic model of Avestimehr et al. [3]. An interesting observation is that even within the signal level dimension that is captured by the deterministic channel model, vector space alignment techniques are used to find DoF characterizations. On the other hand, deterministic models of [3] may not be as useful when multiple antennas are involved. In these cases, such as the  $(1 \times N)^{N+1}$  interference channel explored in [53] (see Appendix B), the necessary insights for GDoF characterizations may be accessible from a more fundamental deterministic setting — the deterministic interference model of El Gamal and Costa [42].

In addition to the evolution of lattice alignment described above, there is another research thread, on computation capacity, which originated as early as 1979 with Korner and Marton’s work in [82], which more recently evolved into a class of compute-and-forward schemes due to Nazer and Gastpar’s work in [100], and which has naturally incorporated lattice alignment principles in its recent applications, such as [103, 106]. The use of structured codes for distributed dirty paper coding to achieve the capacity of the double dirty MAC channel, by Philosof and Zamir in [119], also attributed to a form of lattice alignment, is intriguing as well.

In summary, on the one hand lattice codes along with lattice alignment offer the opportunity to decode and subtract aggregate interference even when interferers are not individually decodable. This indicates a strong advantage for structured codes in interference networks. On the other hand, because of the sensitivity of lattice alignment approaches to channel uncertainty the debate is not entirely over. If interference alignment is not possible with the amount of channel uncertainty that is unavoidable in practice, then the advantage of structured codes may be greatly reduced and random coding approaches may turn out to be more robust. As a remarkable counterpoint, encouraging results for robust lattice alignment schemes in the presence of channel uncertainty are presented by Huang et al. in [60].

Indeed, the robustness of interference alignment to channel uncertainty may be the most critical determinant of its impact on future

wireless networks. The following sections review encouraging results in this direction.

## 4.9 Blind Interference Alignment

Blind interference alignment is introduced in [63] as the idea that even if the actual channel coefficient values are entirely unknown to the transmitters, interference alignment may be possible based on the knowledge of the distinct autocorrelation properties of the channels seen by different receivers. As an example consider the two user MISO BC channel where the transmitter is equipped with two antennas and each receiver has one antenna. With full channel state information at the transmitter (CSIT) it is clear that the transmitter can send a noninterfering signal to each user by sending it along a beam orthogonal to the other users' channel vector, and thus 2 DoF are achievable. Without CSIT, it may seem that no multiplexing of signals is possible and only 1 DoF can be achieved, as indeed has been conjectured in [86]. If the channels are statistically indistinguishable then it is easy to see that indeed the MISO BC is degraded and can achieve only 1 DoF [31]. While block fading channels are precluded in the setting of [86], the assumption that the channel stays constant across a few adjacent time and frequency slots is commonly made in wireless channel models to capture notions like coherence time and coherence bandwidth. The blind interference alignment scheme exploits this assumption and the fact that the users are not statistically equivalent.

On the top left corner of Figure 4.8 is a conceptual depiction of the physical autocorrelation model for the two users where the change of colors depicts decorrelation of channel coefficients across time and frequency. Evidently, User 1 has a smaller coherence bandwidth and a larger coherence time than User 2. Three signal dimensions are chosen as indicated in the autocorrelation graph, so that User 1's channel state changes from the first dimension to the second and then stays constant from the second to the third, while User 2's channel stays constant across the first two dimensions and then changes to a different value. The channel states over the three signaling dimensions are indicated in the top right corner of Figure 4.8. The bottom half of the figure

shows how a total of four symbols, two for each user, are transmitted from the two transmit antennas, over the three chosen signaling dimensions. The signaling dimensions are indicated as time slots for ease of exposition. Since there are four symbols being transmitted over three signal dimensions it is clear that the two undesired symbols must align into one dimension for each receiver to be able to recover its own two desired symbols from the remaining two interference-free dimensions.

With the transmission scheme shown in Figure 4.8, the received symbols at the two receivers are expressed as follows:

$$\begin{aligned} \begin{bmatrix} y^{[1]}(1) \\ y^{[1]}(2) \\ y^{[1]}(3) \end{bmatrix} &= \begin{bmatrix} h_1^{[1]}(1) & h_2^{[1]}(1) & 0 & 0 & 0 & 0 \\ 0 & 0 & h_1^{[1]}(2) & h_2^{[1]}(2) & 0 & 0 \\ 0 & 0 & 0 & 0 & h_1^{[1]}(2) & h_2^{[1]}(2) \end{bmatrix} \\ &\quad \times \begin{bmatrix} x_1^{[1]} \\ x_2^{[1]} \\ x_1^{[1]} + x_2^{[1]} \\ x_1^{[2]} + x_2^{[2]} \\ x_1^{[2]} \\ x_2^{[2]} \end{bmatrix} + \begin{bmatrix} z^{[1]}(1) \\ z^{[1]}(2) \\ z^{[1]}(3) \end{bmatrix} \\ \begin{bmatrix} y^{[2]}(1) \\ y^{[2]}(2) \\ y^{[2]}(3) \end{bmatrix} &= \begin{bmatrix} h_1^{[2]}(1) & h_2^{[2]}(1) & 0 & 0 & 0 & 0 \\ 0 & 0 & h_1^{[2]}(1) & h_2^{[2]}(1) & 0 & 0 \\ 0 & 0 & 0 & 0 & h_1^{[2]}(2) & h_2^{[2]}(2) \end{bmatrix} \\ &\quad \times \begin{bmatrix} x_1^{[1]} \\ x_2^{[1]} \\ x_1^{[1]} + x_2^{[1]} \\ x_1^{[2]} + x_2^{[2]} \\ x_1^{[2]} \\ x_2^{[2]} \end{bmatrix} + \begin{bmatrix} z^{[2]}(1) \\ z^{[2]}(2) \\ z^{[2]}(3) \end{bmatrix}. \end{aligned}$$

During time slot 1 the two symbols intended for User 1 are sent from the two transmit antennas. During time slot 3 the two symbols intended for User 2 are sent similarly. Each receiver has thus acquired one linear combination of its desired symbols and one linear combination of the interfering symbols. The crucial part is time slot 2, where

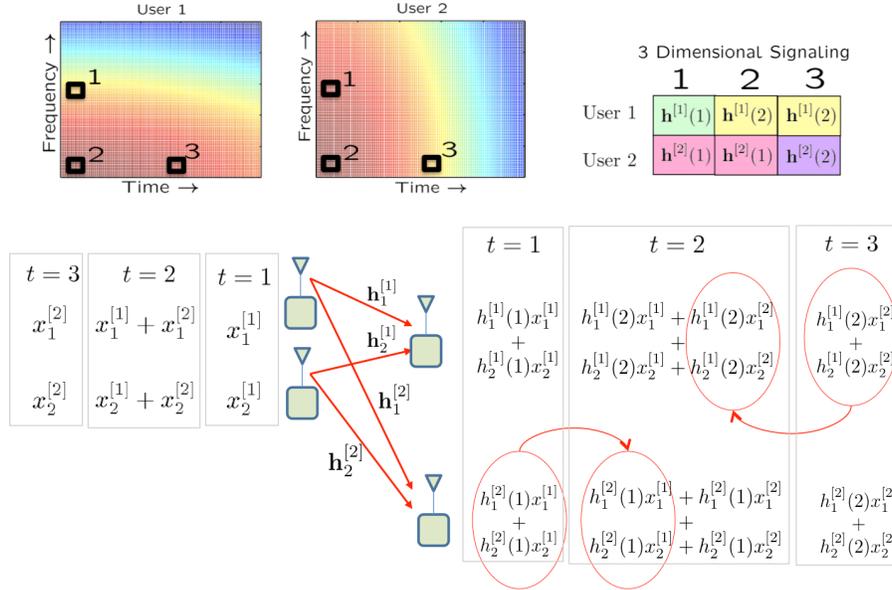


Fig. 4.8 Blind interference alignment scheme [63] for two user MISO BC or the X channel. Symbols  $x_1^{[k]}, x_2^{[k]}$  are desired by User  $k$ .

the transmitted symbols are a superposition of time slots 1 and 3. For User 1 this means that during time slot 2, he receives a different linear combination of his desired symbols as in time slot 1 (because his channel in time slot 2 is different from time slot 1) but the same linear combination of interfering symbols as time slot 3 (because his channel in time slot 2 is the same as in time slot 3). This allows User 1 to subtract the linear combination of interfering symbols that is observed twice and thus he is left with two interference-free different linear combinations of his own desired symbols from which the two symbols can be resolved. Similar arguments can be applied to User 2, by symmetry of construction. Thus, each user is able to eliminate interference and recover his own two desired symbols. Since this occurs over a total of 3 time slots, the DoF achieved is  $4/3$ . Incidentally this is also the DoF outer bound for this setting.

There are several interesting aspects of the blind interference alignment scheme of Figure 4.8. First, note that no data sharing is needed between the two transmit antennas. So, the same scheme can be applied

to the X channel setting as well, where instead of a single transmitter with two antennas, there are two transmitters with a single antenna each. In fact, Figure 4.8 already depicts the two transmit antennas as distributed. The second observation is that in order to cancel interference, the receivers do not need to know the values of the channel coefficients. Of course, one interference is removed, channel knowledge is needed by the receivers in order to perform coherent detection of their own desired signals.

Similar blind interference alignment schemes are also shown in [63] for the two user MIMO interference channel and the  $K$  user SISO interference channel. The two user MIMO interference channel setting is one where User 1 has one transmit and two receive antennas, while User 2 has three transmit and four receive antennas. With no CSIT at the users, the question of interest is that if User 1 achieves 1 DoF, what is the maximum DoF simultaneously achievable by User 2. It is shown that if User 1's desired channel varies faster than the interfering carrying channel, then 1.5 DoF can be achieved by User 2, which requires interference alignment. Since without CSIT, Transmitter 2 cannot zero force its signal at Receiver 1, it must align the 1.5 dimensional signal into 1 dimension, or equivalently, operating over two time slots, align the 3 dimensional transmitted signal into 2 dimensions at Receiver 1.

An alternate setting for blind interference alignment is proposed by Gou et al. in [54] in the form of staggered antenna switching. The main idea is that channel coherence block patterns can be manipulated by changing the antenna characteristics in a pre-determined fashion either through reconfigurable antennas or through antenna switching. For the MISO BC with  $M$  antennas at the transmitter and  $K$  single-antenna receivers capable of switching blindly among  $M$  reconfigurable modes, it is shown that the DoF outer bound of  $\frac{MK}{M+K-1}$  is achievable by staggered mode switching. Note that this is the same as the DoF outer bound of the X channel obtained by breaking the  $M$  transmit antennas into  $M$  single antenna transmitters, and like the example of Figure 4.8, is also achievable with the distributed transmitters. Therefore, interestingly, the X channel does not lose DoF relative to the perfect CSIT setting even when no CSIT is present as long as the receivers have the ability to switch among independent antenna modes. Wang et al. also

extend their results to the MIMO BC setting in [153]. Blind interference alignment through staggered antenna switching is explored for the two user MIMO interference channel by Ke and Wang in [79] and the achieved DoF are shown to be tight with respect to an outer bound obtained earlier in [59, 61, 66]. Benefits of blind interference alignment schemes in cellular settings as well as rate improvements possible through optimal power allocation are investigated by Wang et al. in [156, 157].

It should also be noted that if channel autocorrelations across adjacent time/frequency slots are ignored through the assumption of i.i.d. fading, especially when the fading model is isotropic, then the DoF benefits of blind interference alignment are lost. The DoF results for i.i.d. isotropic fading models are obtained in [59, 148, 166, 167] and in all cases it is shown that the DoF benefits of interference alignment are lost. Since the realistic fading models are somewhere in between the i.i.d. fading models of [59, 148, 166, 167] and the block fading models of [54, 63, 79, 153], the benefits of blind interference alignment in practical settings remain an active research area.

#### 4.10 Retrospective Interference Alignment

Retrospective interference alignment is introduced by Maddah-Ali and Tse in [94] and refers to interference alignment schemes that exploit only delayed CSIT. The delayed CSIT is assumed independent of the current channel state. The general intuition prior to the Maddah-Ali and Tse's work in [94] was that channel knowledge is only useful to the transmitter to the extent that it helps learn the current channel state. It is therefore surprising that even delayed CSIT that is independent of the current channel state is, not only useful, but so useful that it can increase the available DoF. As an illustrative example, consider the MISO BC with two antennas at the transmitter and two users with one receive antenna each. For retrospective interference alignment the channels can (but do not have to) be i.i.d. isotropic. In the absence of delayed CSIT, i.i.d. isotropic fading channels would lose all signal multiplexing benefits and only have 1 DoF, as mentioned in the previous section. However, Maddah-Ali and Tse show that the two user,

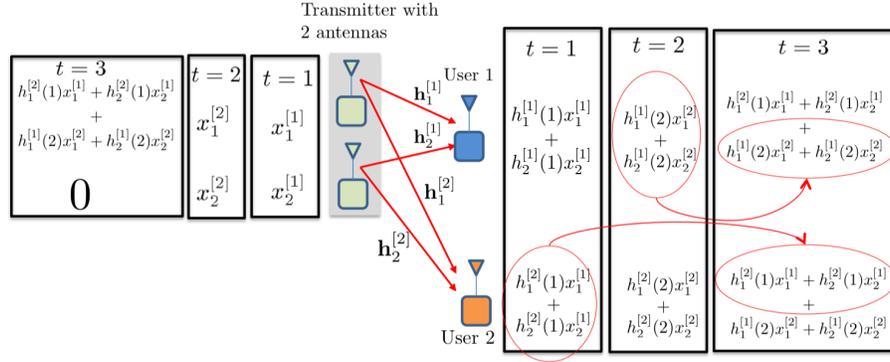


Fig. 4.9 Retrospective interference alignment scheme [94] for two user MISO BC. Symbols  $x_1^{[k]}, x_2^{[k]}$  are desired by User  $k$ .

two transmit antenna MISO BC achieves 4/3 DoF, with only delayed CSIT available at the transmitter. Interestingly, the achievable scheme, shown in Figure 4.9 is very similar to the blind interference alignment scheme illustrated in Figure 4.8 if time slots 2 and 3 are switched in the latter.

In time slot 1, the two symbols for User 1 are sent from the two antennas of the transmitter. User 1 hears a linear combination of desired symbols, while User 2 hears a linear combination of undesired (interfering) symbols. In time slot 2, the same action is performed with the two symbols meant for User 2, so that User 2 hears a linear combination of desired symbols while User 1 hears a linear combination of undesired symbols. No CSIT is used so far. The crucial interference alignment step occurs in time slot 3. Now the transmitter has knowledge of all the channel states in the first two time slots (note that the transmitter does not know the current channel state). Based on this channel knowledge the transmitter is able to reconstruct the linear combinations of undesired symbols seen by each receiver. The transmitter simply adds these linear combinations and sends their superposition *from the same transmit antenna*. In the third time slot, each receiver inverts its channel, so that the channel gain to the transmitting antenna can be normalized to unity. This gives each receiver a superposition of the linear combination of interfering symbols that it has already seen,

and a new linear combination of desired symbols that it has not seen previously. Subtracting the previously seen linear combination of interfering symbols, each receiver has access to two different linear combinations of its desired symbols from which it can resolve the two desired symbols. The scheme is generalized in a nontrivial fashion to the MISO BC with  $K$  transmit antennas and  $K$  single antenna receivers where a total of  $\frac{K}{1+\frac{1}{2}+\frac{1}{3}+\dots+\frac{1}{K}}$  DoF are achieved. Moreover if all receivers have identically distributed channels, then a converse argument based on an outer bound of [59], along with the result that feedback does not increase the capacity of a physically degraded BC, is used to show that this is also the maximum DoF possible with delayed CSIT.

The MISO BC setting of [94] is generalized to the MIMO BC with arbitrary number of antennas at each node by Vaze and Varanasi in [150] under the same assumption of delayed CSIT. Vaze and Varanasi obtain a general outer bound for the  $K$  user setting and show that it is tight for the  $K = 2$  user case, so that for  $K = 2$  the DoF region is given by the bound  $\frac{d_1}{\min(M, N_1 + N_2)} + \frac{d_2}{\min(M, N_2)} \leq 1$  and its counterpart obtained by switching indices 1, 2. Here  $M$  is the number of antennas at the base station, and  $N_k$  is the number of antennas at Receiver  $k$ ,  $k \in \{1, 2\}$ . To appreciate the role of CSIT in the two user MIMO BC, suppose, that  $N_2 \geq N_1$ . The DoF region for this setting is found for perfect CSIT by Weingarten et al. in [160], for no CSIT by Huang et al. in [59], and for delayed CSIT by Vaze and Varanasi in [150]. When  $M \leq N_1$  the DoF region is identical with perfect, delayed or no CSIT. As  $M$  increases so that  $N_1 < M \leq N_2$ , the DoF region with perfect CSIT is bigger than that with no CSIT or delayed CSIT, the last two being identical. So, there is no advantage of delayed CSIT in terms of the DoF region, unless  $M > N_2$ , at which point delayed CSIT leads to a larger DoF region than no CSIT. Lastly, increasing  $M$  beyond  $N_1 + N_2$  does not further enlarge the DoF region whether the setting is with perfect, delayed or no CSIT.

While there are some similarities between the blind interference alignment example of Figure 4.8 and the retrospective interference alignment example of Figure 4.9, there are also very significant fundamental differences. The blind interference alignment scheme of Figure 4.8 is equally applicable to the MISO BC and the X channel,

i.e., it pays no DoF penalty for breaking up the two antenna transmitter into two transmitters with a single antenna each. On the other hand, the retrospective interference alignment scheme of Figure 4.9 *requires* data sharing between the two transmit antennas for the third time slot transmission. Thus it cannot be applied as such to the X channel. For the blind interference alignment scheme, one receiver does not need to know the channel state of the other receiver. However, for retrospective alignment, each receiver needs to know the channel state of the other receiver. This is because the linear combination of desired symbols received over the third time slot is formed with the other users' channel coefficients.

For the setting with *distributed transmitters*, especially when an independent message originates at each transmitter, a distributed retrospective interference alignment scheme is introduced by Maleki et al. in [95]. The distributed retrospective interference alignment scheme is capable of achieving  $8/7$  DoF for the two user X channel, as well as for two interfering MAC channels with only single-antenna nodes, each consisting of two users communicating with a receiver. It is not known, however, if  $8/7$  DoF is also the best achievable DoF for these settings with delayed CSIT. Similarly, distributed retrospective interference alignment is used to achieve  $9/8$  DoF on the three user interference channel with single antenna nodes, but it is not known if this is the maximum value of DoF with delayed CSIT for this channel.

Maleki et al. also provide an example of retrospective interference alignment for the two-user MIMO interference channel. The setting  $(1 \times 2, 1)(3 \times 4, d_2)$  was first highlighted by Huang et al. in [59] where the question of the largest possible value of  $d$  without perfect CSIT was posed. Huang et al. argued that depending on whether or not interference alignment is possible, the maximum feasible value of  $d_2$  could be as high as 1.5 or as low as 1. Here feasibility is intended in the information theoretic sense, and not necessarily restricted to linear schemes. With no CSIT, if all channels are i.i.d. isotropic, Zhu and Guo show in [167] that no interference alignment is possible and the maximum possible value of  $d_2$  is 1. In the absence of CSIT, if the channels are not isotropic, the largest value of  $d_2$  remains open.

With delayed CSIT, Maleki et al. show in [95] that indeed  $d_2 = 1.5$  is achievable through retrospective interference alignment.

The complete DoF region of the two user MIMO IC with delayed CSIT and with arbitrary number of antennas at each node, was obtained by Vaze and Varanasi in [149]. Interestingly, in some cases delayed CSIT achieves the same DoF region as when full CSIT is available, and which is strictly larger than possible with no CSIT (in the isotropic fading case). For example, the DoF point  $(1 \times 2, 1)(4 \times 3, 2)$  is on the DoF region boundary with perfect CSIT, is also achievable with delayed CSIT, but cannot be achieved without CSIT under i.i.d. isotropic fading.

Maleki et al. [95] also consider the possibility of providing delayed channel output feedback to the transmitters instead of delayed CSIT. While neither setting is clearly superior to the other, it is shown that with channel output feedback the X channel can achieve  $4/3$  DoF (i.e., the same as the MISO BC), and the three user interference channel can achieve  $6/5$  DoF, respectively. The DoF gains enabled by delayed channel output feedback are especially remarkable given that with full instantaneous CSIT, it has been shown by Cadambe and Jafar in [28] that feedback does not increase the DoF of fully connected X or interference networks with any number of users.

# 5

---

## Applications of Interference Alignment

---

The idea of interference alignment is a common ingredient of many, if not most, multiuser communication network capacity problems. Given a communication network, a fundamental challenge is to determine what form of interference alignment may lead to the most accessible signal dimensions. While there is no completely systematic or general approach known so far that is applicable to all settings where interference alignment is needed, a focus of this section is to highlight the broad range of problems where the asymptotic interference alignment scheme of [CJ08] is applicable. We expect that by the end of this section, the reader will be able to develop a systematic understanding of the [CJ08] scheme and to apply it even to networks not considered anywhere previously. As discussed previously in Section 4.7 the applications of the [CJ08] scheme can be interpreted either over time-varying/frequency-selective channels to achieve alignment of signal vector spaces through linear beamforming, or over constant channel coefficients (both real or complex) to achieve alignment of rational dimensions in the framework of [93, 99]. As such, we will not distinguish between these settings in the descriptions below. Finally, we will assume global and perfect channel knowledge for all examples in this section.

## 5.1 $K$ User Interference Channel

The  $K$  user interference channel has  $K/2$  DoF in general. The achievable scheme is based on the asymptotic alignment scheme of [CJ08], described previously in Section 4.6, and is illustrated again in Figure 5.1(a). The columns of  $\mathbf{V}$  are the signaling dimensions along which individual streams are sent. These columns are designed according to the [CJ08] construction to be (approximately) invariant to linear transformations  $T_i$  corresponding to the interfering links shown in red in Figure 5.1(a). All transmitters use the same  $\mathbf{V}$ . The invariance of  $\mathbf{V}$  to the cross-channels forces all the interference at each receiver into the same signaling dimensions  $\mathbf{V}$ . The direct channels  $H^{[kk]}$  act as generic linear transformations that cause the desired signals  $H^{[kk]}\mathbf{V}$  to have the minimum overlap possible with the interference space. By choosing the dimension of the overall space to be big enough to contain both desired and interfering dimensions, the overlap between desired and interfering signal dimensions is reduced to zero. Since  $H^{[kk]}$  are full rank transformations, the number of dimensions of desired signals  $H^{[kk]}\mathbf{V}$  is the same as the number of dimensions occupied by interference  $\mathbf{V}$ . Thus, at each receiver, half the total number of signal dimensions is occupied by

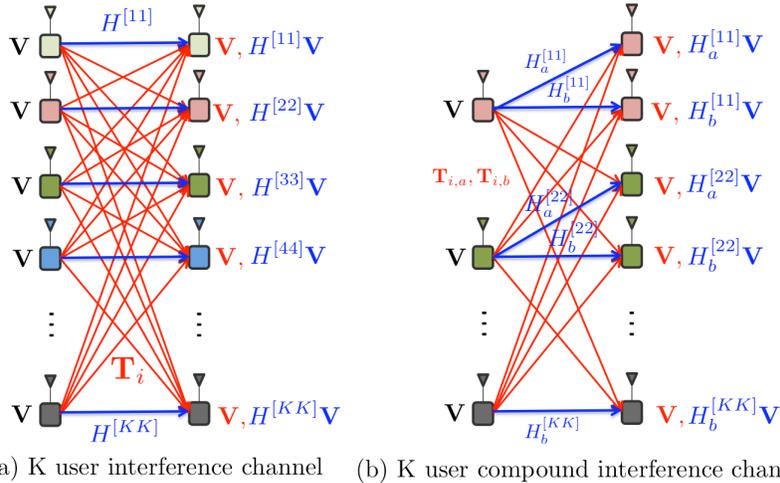


Fig. 5.1 [CJ08] scheme to achieve  $K/2$  DoF on the  $K$  user interference channel, without compound states on the left, and with compound states on the right.

interference, and the remaining half is available to the desired signals free from interference. In other words, everyone gets half the cake.

Going beyond sum-DoF to the entire DoF region, a simple time-sharing argument shows that the DoF region of the  $K$  user interference channel consists of DoF tuples  $(d_1, \dots, d_K)$  that satisfy all the inequalities  $d_i + d_j \leq 1$ ,  $\forall i, j \in \{1, 2, \dots, K\}, i \neq j$ . Moreover, if every node is equipped with  $M$  antennas, the DoF region is simply enlarged by a factor  $M$ . This is shown for  $K = 3$  in [22] and for general  $K$  in [163]. Further generalizations of the  $K$  user interference channel include the setting with general message sets considered in [78] where the [CJ08] scheme is used to find the DoF region for the case where we have  $K$  transmitters, each with an independent message, and  $J$  receivers, each of whom may be interested in arbitrary subsets of the  $K$  messages.

Another illustrative example is the finite state compound setting for the  $K$  user interference channel, as shown in Figure 5.1(b). The compound setting may be used to model channel uncertainty, or be interpreted as multiple multicasts. In a  $S$  state compound channel, each receiver can be in one of  $S$  distinct states, or equivalently, there is a group of  $S$  receivers that are all interested in the same message. Thus, instead of  $K$  receivers, each interested in one message, we now have  $K$  groups of receivers, with each group interested in one message, such that all receivers in that group must decode the desired message. Clearly the compound setting makes the interference alignment task more challenging because of the increase in the number of alignment constraints. The conventional wisdom for compound channels was that as the number of states increases, the increasing number of signal zero-forcing and alignment constraints, will lead to a collapse of the DoF. However, as shown in Figure 5.1(b) the [CJ08] construction applies directly to the compound channel setting as well with no loss of DoF, i.e., regardless of the number of states (assumed finite), each user achieves half the cake.

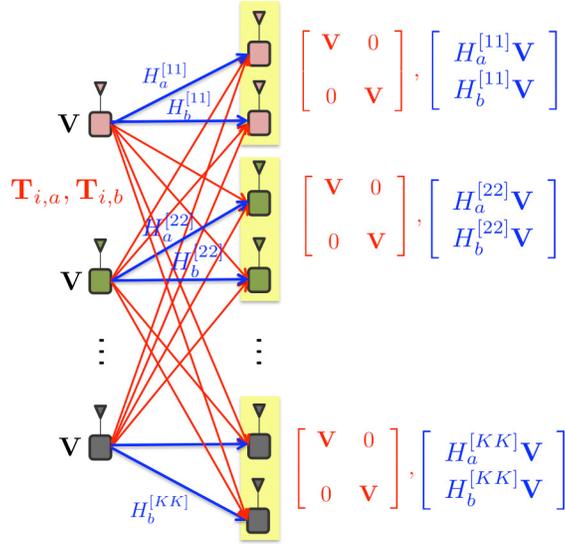
*Open Problem:* While the DoF of the  $K$  user interference channel are robust to the finite state compound setting, a challenging open problem is to determine the DoF when the number of states is countably infinite, or worse, uncountable.

## 5.2 $K$ User $M \times N$ MIMO Interference Channels

A  $K$  User  $M \times N$  MIMO interference channel refers to the  $K$  user interference channel where every transmitter has  $N$  antennas and every receiver has  $M$  antennas. Going from the DoF characterization of the SISO ( $M = N = 1$ )  $K$  user interference channel to the MIMO  $K$  user interference channel is straightforward when all nodes have the same number of antennas,  $M = N$ . In that case, simply by splitting each node into  $M$  separate nodes, we transform the  $K$  user  $M \times M$  MIMO interference channel into the  $KM$  user  $1 \times 1$  interference channel, for which we already know that  $MK/2$  DoF are achievable. Since this is also the outer bound for the  $K$  user  $M \times M$  MIMO interference channel, the DoF characterization is complete.

The characterization of the DoF of MIMO interference channels is more challenging when all the nodes do not have the same number of antennas. For example, consider the  $K$  user  $M \times 1$  SIMO interference channel where each receiver has  $M$  antennas while each transmitter has only one antenna. Gou and Jafar have shown in [50] that this channel has  $\frac{MK}{M+1}$  DoF in total, i.e.,  $\frac{M}{M+1}$  DoF per user, when  $K > M$ . Since in the absence of interference each user can achieve only 1 DoF, each user achieves  $\frac{M}{M+1}$  fraction of the cake. As a sanity check note that in the SISO setting ( $M = 1$ ), this fraction reduces to  $1/2$ , as expected.

To see how this is done, consider the  $M = 2$  case, as shown in Figure 5.2. Let us first consider the interference seen by any receiver, say Receiver 1. As before, all transmitters use the same signal space  $\mathbf{V}$ . Note that because the receiver has twice as many antennas as the transmitter, the size of the signal space seen by Receiver 1 is twice as big as the transmit space accessible to any transmitter individually. As we know from the property of generic subspaces, they do not overlap if they do not have to, i.e., if the overall space is big enough to accommodate both of them. Therefore, under generic channel conditions the transmit spaces from any two transmitters will not overlap. In other words, interference from any two transmitters cannot be aligned with each other as seen from Receiver 1. This may seem at first a challenging problem, because so far all the interference alignment schemes discussed in Section 4 have relied on one-to-one interference alignment, i.e., an

Fig. 5.2 [CJ08] scheme used to achieve the DoF of the  $K$  user SIMO interference channel.

interference carrying signal vector from one transmitter is aligned with an interference carrying signal vector from another transmitter. In the SIMO interference channel, on the other hand, we see that the interference from no two transmitters can align with each other.

The key to this puzzle is the idea of many-to-many alignment, or the alignment of subspaces instead of alignment of vectors themselves. For instance, in the introductory example (2.4)–(2.6) the interfering vectors  $\mathbf{H}_{\star 2}, \mathbf{H}_{\star 3}, \mathbf{H}_{\star 4}, \mathbf{H}_{\star 5}$  do not possess any one-to-one alignment. The vectors are pairwise linearly independent. However, the subspace spanned by vectors  $\mathbf{H}_{\star 4}, \mathbf{H}_{\star 5}$  aligns perfectly with the subspace spanned by vectors  $\mathbf{H}_{\star 2}, \mathbf{H}_{\star 3}$ . Applied to the SIMO interference channel setting of Figure 5.2, what this means is that while the interference signals from any two transmitters cannot be aligned among themselves, the interference from any *three* transmitters can be aligned within the space spanned by the interference from any two of them. In general, i.e., in the  $K$  user  $M \times 1$  SIMO interference channel, the interference signals from any  $M$  transmitters cannot be aligned among themselves, but the interference from any  $M + 1$  transmitters can be aligned within the space spanned by the interference from any  $M$  of them.

Returning to the SIMO interference channel of Figure 5.2, since all transmitters use the signal space  $\mathbf{V}$  and the signals from any two transmitters cannot overlap, the interference space from transmitters 2 and 3 must already occupy  $2|\mathbf{V}|$  dimensions at Receiver 1. Now the key is to align the interference from all the remaining undesired transmitters into the  $2|\mathbf{V}|$  dimensional interference space occupied by the signals from transmitters 2 and 3. We proceed as follows.

The interference space needs to be twice as large as the transmit signal space. We will set the interference space at each receiver to  $\mathbf{V} \times \mathbf{V}$ , i.e., the column span of the matrix

$$\begin{bmatrix} \mathbf{V} & \mathbf{0} \\ \mathbf{0} & \mathbf{V} \end{bmatrix}.$$

Indeed within this interference space, all interference can be aligned. For example, consider a transmitter with interference-carrying matrices  $\mathbf{T}_{i,a}, \mathbf{T}_{i,b}$  to the two antennas of Receiver 1. The alignment we seek is between the actual interference space, i.e., the column-span of the matrix of received interference vectors and the designated interference space, i.e.,  $\mathbf{V} \times \mathbf{V}$ :

$$\begin{aligned} \text{span} \begin{bmatrix} \mathbf{T}_{i,a} \mathbf{V} \\ \mathbf{T}_{i,b} \mathbf{V} \end{bmatrix} &\stackrel{\subset}{\approx} \text{span} \begin{bmatrix} \mathbf{V} & \mathbf{0} \\ \mathbf{0} & \mathbf{V} \end{bmatrix}, \\ &\Leftrightarrow \mathbf{V} \approx \mathbf{T}_{i,a} \mathbf{V} \approx \mathbf{T}_{i,b} \mathbf{V} \end{aligned}$$

where the symbols  $\stackrel{\subset}{\approx}, \approx$  are meant to highlight that the inclusion and equivalence will only be achieved in the approximate and asymptotic sense of the [CJ08] scheme.

This is quite remarkable — we have reduced a many-to-many alignment problem, i.e., a subspace alignment problem, into a one-to-one alignment problem that we know how to solve (in an asymptotic sense). Collecting similar alignment conditions from all transmitters to all receivers, once again we find exactly the same alignment problem that the [CJ08] scheme is designed to solve. Thus, the problem is solved by constructing  $\mathbf{V}$  such that

$$\mathbf{V} \approx \mathbf{T}_{1,a} \mathbf{V} \cup \mathbf{T}_{1,b} \mathbf{V} \cup \cdots \cup \mathbf{T}_{N,a} \mathbf{V} \cup \mathbf{T}_{N,b} \mathbf{V}.$$

With this scheme  $\approx 2|\mathbf{V}|$  dimensions are occupied by the interference at each receiver. The desired signal occupies  $|\mathbf{V}|$  dimensions. The overall

signal space at the receiver needs to be  $\approx 3|\mathbf{V}|$  dimensions to keep desired signal separate from interference. Thus, the desired signal at each receiver occupies  $1/3$  of the total dimensions available to the receiver. Since each receiver has two antennas, i.e., two spatial dimensions per channel use, the accessible DoF per user is  $2/3$  per channel use.

The [CJ08] scheme applied to the SIMO interference channel as summarized above, can also be directly translated to the reciprocal setting — the MISO interference channel by the duality property established in [23, 48]. DoF for the  $K$  user interference channel with  $M$  antennas at each transmitter and  $N$  antennas at each receiver are found in [50] when the ratio  $R = \max(M, N)/\min(M, N)$  is an integer, and the result is shown in Figure 5.3. Interestingly, there is no DoF penalty from interference while the number of users  $K \leq R$ , i.e., zero forcing suffices and everyone gets  $\min(M, N)$  DoF, i.e., everyone gets all the cake. On the other hand, when  $K \geq R + 1$ , the DoF per user are again not limited by the number of users, and everyone gets  $\min(M, N) \frac{R}{R+1}$  DoF because of interference alignment. However, there is a clear transition from the zero-forcing regime to the interference alignment regime as we go from  $K = R$  to  $K = R + 1$ . Here the total DoF does not increase even though one user is added. For instance, the

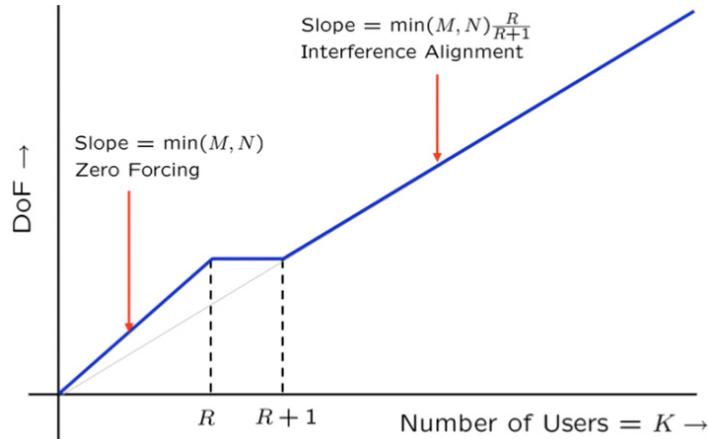


Fig. 5.3 DoF of the  $K$  user  $M \times N$  MIMO Interference Channel when  $R = \frac{\max(M, N)}{\min(M, N)}$  is an integer [50].

two user interference channel has 1 DoF, as does the one user interference channel (i.e., just the point to point channel with no interferer). The  $K = R + 1$  setting is indeed special as it is the highest number of users for which capacity characterizations within an SNR independent bounded gap (i.e., capacity within  $\mathcal{O}(1)$ ) are available for a given integer  $R$  value [53] for almost all channel realizations.

Finally, Ghasemi et al. have shown in [44] that the DoF of the  $K$  user  $M \times N$  interference channel equal  $\frac{R}{R+1}$  per user even when  $R$  is not an integer, provided that the number of users  $K > \frac{M+N}{\gcd(M,N)}$ , where  $\gcd$  refers to the greatest common divisor function. The [CJ08] alignment scheme applied within the rational dimensions framework for constant channels, also achievable equivalently through linear beamforming over time-varying channels, is the key to this result as well.

*Open Problem:* It is evident that multiple antennas create additional signal dimensions, as does the scaling of power. For instance increasing power from  $P$  to  $P^2$  doubles the signal level dimension, while doubling the number of antennas doubles the number of spatial dimensions. To what extent can one dimension be traded off against another, is an interesting open problem of fundamental interest. Another fundamental problem is the issue of scaling of DoF with the spatial dimension. With full CSI and generic channels, is it true (for every network topology) that scaling the number of antennas at each node by  $\alpha$  will increase the total DoF of the network by the same  $\alpha$  factor? While this is the case for all DoF results found so far, establishing this property in general would be a useful step in understanding the similarity/difference between time/frequency and space dimensions.

### 5.3 Cellular Networks

The DoF of cellular networks are characterized by Suh and Tse in [139]. In this setting we have  $K$  cells, each containing a base station and  $M$  users that are served by that base station. All nodes are equipped with a single antenna. The DoF result for this setting is that each cell is able to access a fraction  $\frac{M}{M+1}$  of the cake. Thus, for large number of users, the DoF per cell approach the interference-free setting. While Suh and Tse propose an interference alignment scheme that they call

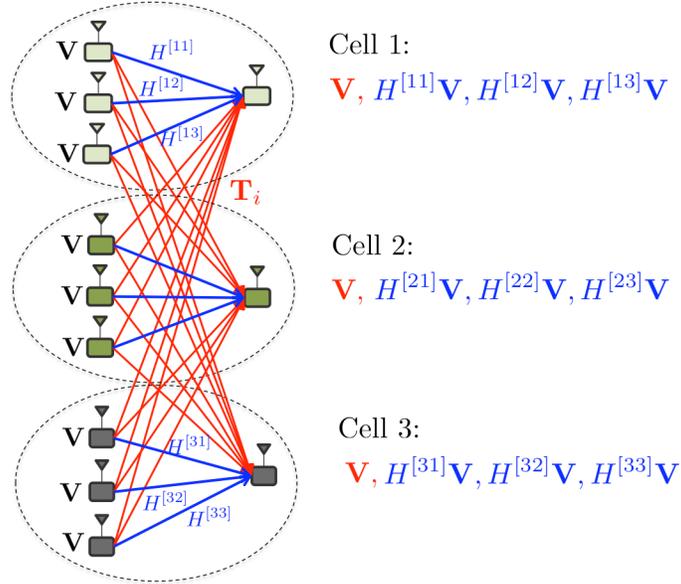


Fig. 5.4 [CJ08] scheme applied to a cellular uplink network with three cells and three user in each cell to achieve  $3/4$  DoF per cell.

subspace alignment in [139], we will continue the theme of this section by presenting a solution based on the [CJ08] scheme.

As an example, consider uplink communication in a network of three cells, with three users in each cell, illustrated in Figure 5.4. As usual, each transmitter uses the same signal space  $\mathbf{V}$ . The interference experienced by a base station, e.g., Base Station 1, consists of terms like  $\mathbf{T}_i\mathbf{V}$ . Collecting all the interference terms together, the [CJ08] scheme ensures

$$\mathbf{V} \approx \mathbf{T}_1\mathbf{V} \approx \dots \approx \mathbf{T}_N\mathbf{V}.$$

With this alignment, the interference occupies  $|\mathbf{V}|$  dimensions at each base station. The  $M$  desired signals from desired users must be mutually separable and therefore need to occupy  $M|\mathbf{V}|$  dimensions at their desired base station. The desired and interfering signals do not overlap at the base station receiver when the overall signal space is large enough to accommodate them, i.e., it spans  $\approx (M + 1)|\mathbf{V}|$  dimensions and all the direct channels (shown in blue in Figure 5.4) are generic.

Thus, each user achieves a total of  $\frac{|\mathbf{V}|}{(M+1)|\mathbf{V}|} = \frac{1}{M+1}$  DoF and each cell achieves a total of  $\frac{M}{M+1}$  DoF.

The MISO downlink cellular setting where base stations equipped with  $M$  antennas transmit to  $N$  single-antenna users in each cell is explored in [44] and it is shown that  $\frac{MN}{M+N}$  DoF are achievable. Interestingly, here also the DoF per cell approach the interference-free setting, i.e., each cell achieves close to  $M$  DoF as the number of users per cell becomes large. The MIMO cellular downlink setting with two cells, each cell having  $K$  users with  $K$  antennas at each user and each cell having a base station equipped with  $K + 1$  antennas is shown to achieve  $2K$  DoF in [133]. At this point we take a brief detour to highlight an interesting observation from the last two sections.

*Observation: DoF benefits of colocated antennas disappear with increasing number of alignment constraints*

Figure 5.5 shows four different networks that each have the same DoF. First, Figure 5.5(a) is the  $K$  user SIMO interference channel where each receiver has  $M$  antennas. This setting is reviewed in Section 5.2 and the [CJ08] solution is illustrated in Figure 5.2. If  $K > M$ , then as illustrated in Figure 5.3 this network has a total of  $\frac{MK}{M+1}$  DoF. The reciprocal network is the  $K$  user MISO interference channel shown in Figure 5.5(b) where each transmitter has  $K$  antennas, and it also has the same DoF. Now, consider the interfering MAC channel setting shown in Figure 5.5 where  $K$  cellular uplinks interfere with each other and each cell has  $M$  users. This setting is considered in Section 5.3 and a solution based on the [CJ08] scheme is also illustrated in Figure 5.4. Here also the total DoF =  $\frac{MK}{M+1}$ . The last setting shown in Figure 5.5(d) consists of  $K$  interfering downlinks with  $M$  users in each cell. Since this is the reciprocal network of the interfering uplinks network shown in Figure 5.5(c) it is not surprising that this network also has  $\frac{MK}{M+1}$  DoF.

Now, let us compare the networks of Figures 5.5(a) and 5.5(d), i.e., the SIMO interference channel and the interfering downlink channels. Clearly, the interfering downlink channels represent a more restricted network setting than the SIMO interference channel. In particular, the SIMO interference channel becomes the interfering downlinks channel if no cooperation is allowed between the  $M$  receive antennas of each

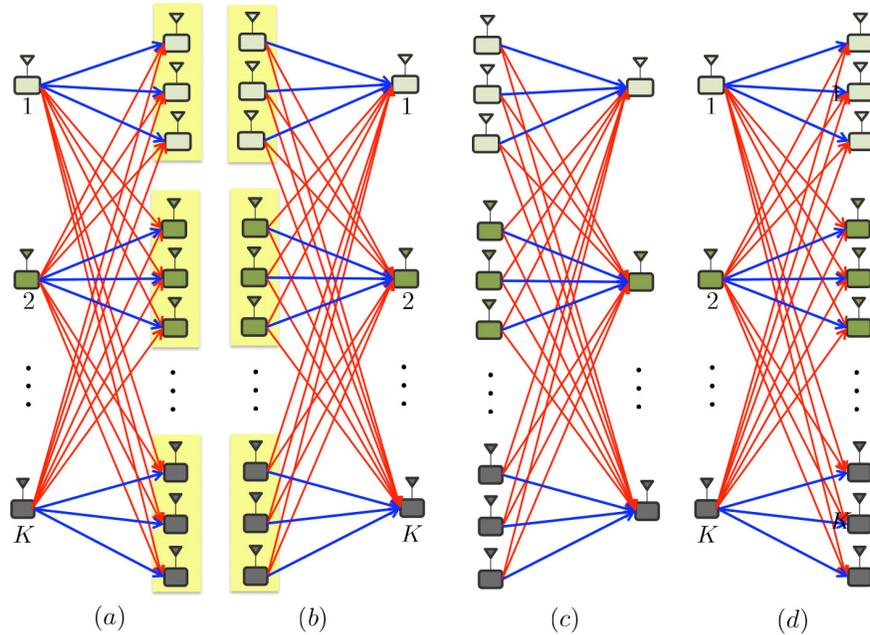


Fig. 5.5 Networks that have the same  $DoF = \frac{M}{M+1}K$  when  $K > M$ , i.e., when the number of alignment constraints exceeds a threshold — (a) The  $K$  user SIMO Interference Channel with  $M$  antennas at each receiver, (b) The  $K$  user MISO Interference Channel with  $M$  antennas at each transmitter (c) Interfering MAC Channels from  $K$  cellular uplinks with  $M$  users in each cell, (d) Interfering BC Channels from  $K$  cellular downlinks with  $M$  users in each cell.

receiver. In general, disallowing cooperation between receive antennas would lead to a loss of DoF. However, as we know these two networks have the same DoF when the condition  $K > M$  is satisfied.<sup>1</sup> Similarly, the network of Figure 5.5(c) is obtained from the network of Figure 5.5(b) by disallowing cooperation among transmit antennas, which also does not lead to a loss of DoF when  $K > M$ . This is a common theme that will reappear in Section 5.5, especially highlighted in Figure 5.7 where again, the price paid for too many alignment constraints will be the loss of the DoF benefits of multiple antennas.

Lastly we point out that the [CJ08] scheme illustrated as a solution to the  $K$  user SIMO interference channel problem in Figure 5.2

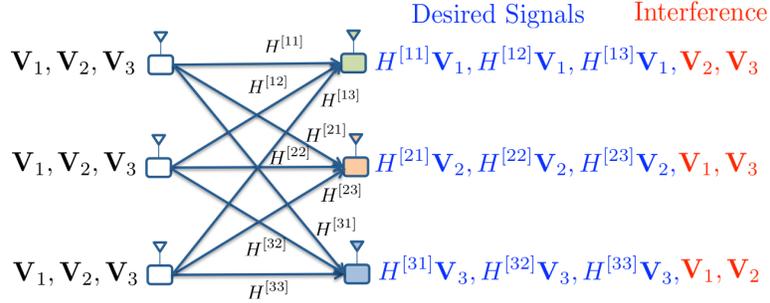
<sup>1</sup>For  $K \leq M$  indeed the SIMO channel benefits from joint processing among the  $M$  receive antennas as it allows zero forcing at receivers and therefore achieves higher DoF than possible in the interfering downlinks setting.

does exploit the joint processing among the  $M$  receive antennas. An alternate scheme that does not require this joint processing can be constructed from the reciprocal version of the [CJ08] solution for the interfering uplinks problem illustrated in Figure 5.4. Thus, there are two distinct solutions possible for the same problem, that are equivalent in DoF, but upon closer inspection will show that much larger symbol extensions are needed in the case without joint processing among the  $M$  receive antennas. In fact the size of symbol extensions is larger by a factor of about  $M$  for the case with separate receive antennas relative to the case where the  $M$  receive antennas can process their signals jointly, i.e., Figure 5.5(a) versus Figure 5.5(c). The distributed solution, i.e., requiring no cooperation among antennas is presented explicitly by Ghasemi et al. for the constant channels setting, i.e., using the rational dimensions framework, in [44].

*Open Problem:* While the  $K$  user SIMO interference channel is equivalent to the interfering BC channels setting in DoF, the two may behave distinctly in GDoF. Interestingly, the GDoF characterization of the  $N + 1$  user,  $1 \times N$  SIMO interference channel is known from the work of Gou and Jafar in [53]. Thus, finding the GDoF of, e.g., the three interfering SISO downlinks problem with 2 users in each cell, will allow direct comparisons against the GDoF of the three user  $1 \times 2$  SIMO interference channel and shed light on the benefits of joint processing of signals at the receivers.

## 5.4 X Networks

Cadambe and Jafar determine the DoF of  $M \times N$  X networks in [23]. An  $M \times N$  X network consists of  $M$  single antenna transmitters,  $N$  single antenna receivers and  $M \times N$  independent messages, one from each transmitter to each receiver. Cadambe and Jafar show that a total of  $\frac{MN}{M+N-1}$  DoF are achieved in this setting with time-varying channels. The result is based on the [CJ08] IA scheme and is translated into the rational dimensions framework to achieve the same DoF with constant channels by Motahari et al. in [99]. The application of the [CJ08] IA scheme in this setting is described below on the  $3 \times 3$  X network shown in Figure 5.6.

Fig. 5.6 Interference alignment on the  $3 \times 3$  X channel.

Consider the three independent messages intended for Receiver 1, one from each transmitter. In order to send these messages, each transmitter selects the same signal space  $\mathbf{V}_1$ . The [CJ08] IA scheme ensures that these signals align into the same interference space  $\mathbf{V}_1$  at Receivers 2, 3 where they constitute interference. The subspace constraints for this alignment involve only the channels to Receivers 2, 3, i.e., the channels to Receiver 1 are not involved. This independence of the desired channels, their generic nature, and the choice of overall signal space big enough to accommodate generic subspaces without forcing overlaps between them, is sufficient to ensure that the desired signal spaces at Receiver 1 do not overlap. Similarly, the three messages for Receiver 2, one from each transmitter, are sent along the same signal space  $\mathbf{V}_2$  and align into that same space at Receivers 1, 3 where they constitute interference, while remaining distinct at Receiver 2, where they are desired. And finally, the three messages for Receiver 3 are sent along  $\mathbf{V}_3$ , align along  $\mathbf{V}_3$  at Receivers 1, 2 and are separable at Receiver 3. The size of the signal spaces are chosen to be equal, i.e.,  $|\mathbf{V}_1| = |\mathbf{V}_2| = |\mathbf{V}_3| = \Delta$ . The size of the total signal space (the number of time/frequency slots) is chosen to be  $\approx 5\Delta$  to accommodate all the signals in general position without any unnecessary overlaps beyond those forced by the [CJ08] IA scheme. The scheme applies similarly to  $X$  networks with arbitrary number of source and destination nodes.

*Open Problem:* The DoF characterization of a MIMO  $S \times D$  X network, when both  $S$  and  $D$  are large (at least greater than 2), and all nodes are equipped with  $M > 1$  antennas, and channel conditions

are generic, is an interesting open problem. This is because for large  $X$  networks the multiplicity of alignment constraints evidently requires the use of the [CJ08] scheme, which does not fully exploit joint processing power of multiple antennas, except by breaking them into separate nodes. Surprisingly, in all settings where the DoF are known so far, the MIMO case with large number of alignment constraints lose all DoF benefits of joint processing at the antennas.

### 5.5 Compound MISO BC Channel

The study of the DoF of the compound MISO BC was initiated by Weingarten et al. in [159], who found the DoF outer bounds as well as nontrivial achievable results based on interference alignment. An example of a compound MISO BC is shown on the left side of Figure 5.7, where a two-antenna transmitter wishes to send three independent messages to three separate groups of users. Each group consists of two users who want the same message. The users within a group may also be seen as different possible states of the same user, to model channel uncertainty at the transmitter. Clearly the presence of multiple users within each group, or multiple potential states for each user, places greater demands on the decodability of the message, as each user within a group is now required to be able to decode the message. It was originally thought that as the number of users within each group (or the

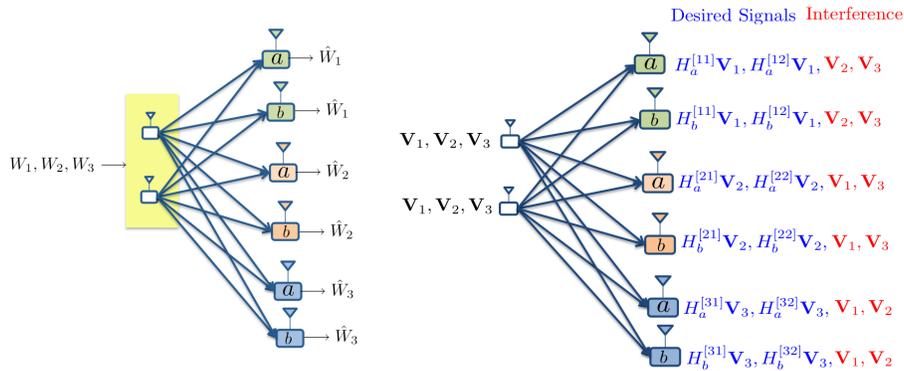


Fig. 5.7 On the left, three user  $2 \times 1$  compound MISO BC with two states for each receiver. On the right,  $2 \times 3$  compound X channel with two states for each receiver.

number of states) increases, the multiplexing gains will be lost and DoF will collapse to 1. However, it was shown by Gou et al. [52] using the [CJ08] scheme that the DoF do not collapse regardless of the number of states for each user. Gou et al. applied the [CJ08] scheme for the time-varying (real or complex) setting and showed how to translate it into the constant (real) channels setting using the rational dimensions framework [99]. The result was extended to *complex* constant channels by Maddah-Ali in [93] based on a fundamental extension of the rational dimensions framework itself to include complex channels.

The key to the DoF of the compound MISO BC when each user's channel can be in one of a large number of states is the underlying X channel obtained by separating the transmit antennas into different transmitters. As shown on the right side of Figure 5.7, the compound BC with no cooperation among transmit antennas, becomes equivalent to a compound X channel. Given this observation, and the ability of the [CJ08] scheme to asymptotically satisfy an arbitrarily large number of alignment constraints with no DoF penalty, it is evident that the compound X channel has the same DoF regardless of the number of states. Thus, a compound X channel with two transmitters and three receivers has  $\frac{2 \times 3}{2+3-1} = 1.5$  DoF *regardless of the number of states of each user*, provided the states are generic and the number of states is finite. Combined with the outer bounding techniques of Weingarten et al. [159], this proved that the DoF of the compound MISO BC do not collapse to unity, but they do collapse to the DoF of the underlying X channel obtained by separating the transmit antennas. Thus, in general, a compound MISO BC with  $M$  transmit antennas and  $K$  single-antenna users, each of which can be in  $J$  different states, has exactly  $\frac{MK}{M+K-1}$  DoF regardless of  $J$ , if  $J$  is large enough — specifically, if  $J \geq M$ .

*Open Problem:* An interesting open problem here is to determine if the DoF equivalence between the compound MISO BC and the compound X channel extends to the GDoF setting. We expect that the benefits of joint processing among antennas will be visible in the GDoF setting and the nature of loss due to distributed processing could provide interesting new insights.

## 5.6 Network Coding — Multiple Unicasts

Along with the recent progress on wireless interference networks, a related success story has been the idea of network coding for communication over wired networks. In spite of the vast differences in the communication problem formulations in the wired and wireless settings, it is interesting to note that these two threads are increasingly converging around the same set of basic principles. For instance, just as recent developments in wireless networks argue against orthogonalization of medium access, the idea of network coding advocates the mixing of packets instead of traditional routing approaches that may be seen as orthogonalizing the flows in a network. While much of the initial success in network coding has been in the context of multicast networks (where all information is desired at all receivers, i.e., there is no undesired signal/interference) it is also becoming clear that some of the most important open problems in network coding, such as the multiple unicast problem, could significantly benefit from the tools developed for wireless interference networks — in particular, from the idea of interference alignment. Once the flows are mixed the network coding problem starts bearing resemblance to the wireless interference channel, as each destination must separate the desired symbols from the undesired symbols based on a set of received linear equations. With linear network coding the end-to-end network allows a channel matrix representation, much like the wireless interference channel and many of the insights from interference alignment schemes may be expected to carry over.

Let us consider the graph shown in Figure 5.8 and three unicast sessions: sources 1, 2, 3 send one symbol,  $x, y, z$  respectively, to their respective destinations. First, we note that the min-cut between any corresponding source–destination pair is 2, and thus, each session could send at rate equal to two symbols per channel-use, if it were the only one routed over this network. However, since all desired flows must pass through only two edges (the ones at the top and the bottom of the figure), routing cannot achieve a sum-rate of more than two symbols per channel-use. On the other hand, when interference align-

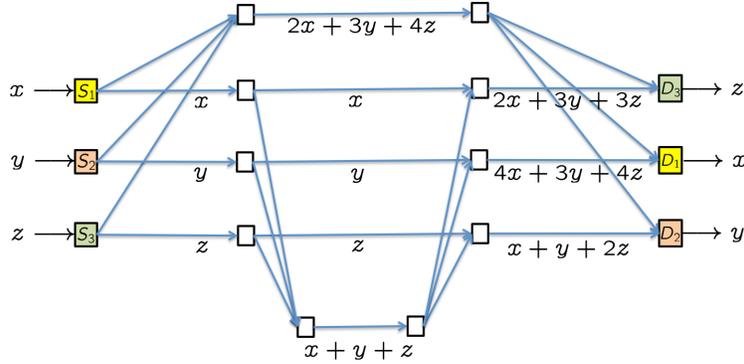


Fig. 5.8 Interference alignment for three Unicast Sessions [120]. Each link has the same capacity, which is equal to the information carried by each of the symbols  $x, y, z$ .

ment is used in conjunction with network coding at intermediate nodes, each session can simultaneously achieve half its interference-free rate. Figure 5.8 shows one way to achieve that goal. Each destination receives two equations in the three variables  $x, y, z$ . Even though there are more variables than equations, alignment allows these two equations to be solved to recover the desired variable. For example, destination 1, subtracts  $2x + 3y + 4z$  from  $4x + 3y + 4z$  to simultaneously eliminate all interference and recover the value of  $x$ . Furthermore, the sum rate of three symbols per channel-use, achieved by interference alignment and network coding can be shown to be capacity optimal for this example. As explained earlier, routing by itself can only achieve a maximum sum-rate of two symbols per channel-use. Random linear network coding (by itself, without alignment) fails because the number of equations is smaller than the number of variables. A careful inspection shows that breaking down the network into butterfly structures does not suffice either. In fact, alignment is necessary to achieve the optimal data rates for this network.

Recent work by Das et al. in [36, 120] explores the use of the [CJ08] alignment scheme to find systematic interference alignment solutions to arbitrary multiple unicast problems. Applications of the [CJ08] scheme to three-source three-destination unicast networks *with delays* are studied by Bavirisetti et al. in [9], where a transform approach is used to convert the network with delays into parallel instantaneous networks.

The approach taken in these works is to identify conditions when half the min-cut is achievable for each flow, based on the alignment guarantees provided by the [CJ08] IA scheme. The principal challenge is the existence of spatial dependencies among channels as multiple flows pass through the same set of intermediate nodes. The spatial dependencies affect the assumption of *generic* direct channels, needed in the [CJ08] scheme for desired signals to remain distinguishable from interference. This work is still in its early stages and the understanding of these spatial dependencies is so far limited to the three unicast setting, where analogies are drawn to the three user interference channel.

Another issue that comes up in translating interference alignment schemes from wireless networks to network coding applications is that the former setting operates over continuous signals whereas the latter is invariably defined over finite fields. In particular the notion of generic channels is often invoked in the wireless setting, to mean that since the channels are drawn from continuous distributions, they will not satisfy any more linear relationships than they must. The corresponding statement over finite fields comes from the Schwartz-Zippel lemma, which implies that the probability that a nonzero polynomial will have roots at randomly selected test points can be made arbitrarily small by operating over a sufficiently large field size. Since linear independence relationships can be formulated equivalently as the nonvanishing property of determinant polynomials, the generic properties of random variables drawn from continuous distributions translate to the random variables drawn from *sufficiently large* finite fields.

While there exist examples [120] where the interference alignment benefits from wireless networks can be directly adapted into the corresponding wired network setting, such a translation cannot be taken for granted in general due to several very fundamental differences in the nature of wired and wireless networks. For instance, consider a wireless  $X$  network where  $S$  sources wish to communicate with  $D$  destinations over wireless channels, and each source has an independent message for each destination for a total of  $S \times D$  messages. This setting provides the maximum DoF that can be obtained between  $S$  sources and  $D$  destination nodes (among all possible choices of message sets) without genie-aided data sharing among the nodes. Assuming generic channels, the DoF

for this setting are shown to be equal to  $\frac{SD}{S+D-1}$  by Cadambe and Jafar in [23], which can only be obtained through interference alignment and is significantly smaller than the network min-cut  $\min(S, D)$ . On the other hand, consider the corresponding network coding problem with  $S$  sources,  $D$  destinations, and  $S \times D$  independent messages, one from each source to each destination. Surprisingly, regardless of the network topology, the sum-capacity of this network is equal to the network min-cut and does not require any interference alignment. In fact, it does not require any network coding either, as simple routing suffices to achieve the min-cut in this setting. To see this, suppose we allow all  $S$  sources to cooperate, and all  $D$  destinations to cooperate, by collectively processing their signals. This creates a single flow problem for which capacity is known to be the min-cut rate, and is achieved with routing with assignment of flows according to the Ford–Fulkerson algorithm. Since routing does not require any mixing of packets, the routing solution does not require any cooperation across source or destination nodes. Therefore, the same min-cut sum-rate is achievable in the original network without cooperation among nodes as well. Thus, the  $X$  network setting, one of the most important settings to highlight the benefits of interference alignment, is trivial in the network-coding context.

*Open Problem:* The general question of using interference alignment techniques and emerging insights from the study of wireless networks to provide systematic guarantees on achievable rates of multiple unicast network coding problems remains a very interesting and open research avenue.

There is a class of network coding problems where interference alignment has been shown to be essential, and the [CJ08] scheme has been applied most successfully to find the optimal network coding solution. This is the so-called distributed storage exact repair problem. Next, we summarize the distributed storage repair problem and the solution based on the [CJ08] scheme.

## 5.7 Distributed Storage Exact Repair Problem

A detailed description of the distributed storage problem is beyond the scope of this monograph. We refer the reader to an excellent survey by

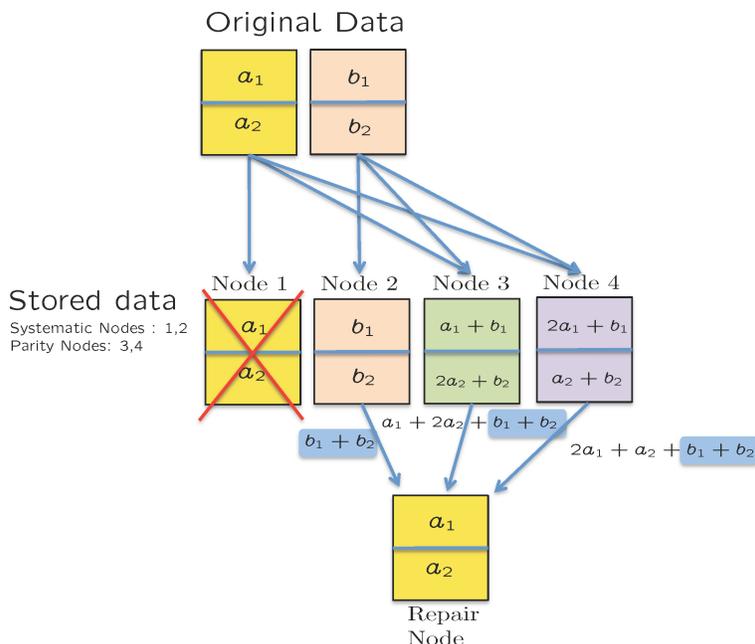


Fig. 5.9 Distributed Storage Exact Repair Problem for a (2,4) MDS code.

Dimakis et al. [37]. For our purpose, we introduce the distributed storage exact repair problem through the example illustrated in Figure 5.9.

Consider a file, conceptually comprised of four data symbols  $(a_1, a_2, b_1, b_2)$  that is to be stored in four distributed storage nodes, each capable of storing two symbols. Nodes 1 and 2 are systematic nodes and store data symbols uncoded, while Nodes 3 and 4 may be regarded as parity nodes and store linear combinations of the symbols. The code is a (2,4) MDS code, i.e., all the information can be recovered from any two storage nodes (so that failure of any two nodes does not result in irreparable information loss), and also the total storage capacity of the nodes that need to survive is equal to the actual data. The latter is also referred to as the minimum storage property. The exact repair problem is concerned with the scenario where one of the storage nodes, e.g., node 1 fails. A new node — the repair node — is required to repair the network by replacing the node with its exact replica. The repair node is capable of contacting some of the surviving storage nodes. However, it

must do so while downloading only the minimum information possible from the nodes that it is able to contact. For this example, the repair node is able to contact all surviving nodes, i.e., nodes 2,3,4. The minimum information to be downloaded by the repair node, also known as the minimum repair bandwidth, is found from an information theoretic min-cut argument to be equal to one symbol from each surviving storage node, for a total of three symbols. Now we note that there are a total of four unknown symbols  $a_1, a_2, b_1, b_2$ . The repair node can only download three symbols, which provide it with three equations. With generic equations, as mentioned earlier, it is not possible to recover the two desired symbols because the number of equations is less than the number of variables. However, with interference alignment, these equations can be cleverly crafted so that the two undesired variables align into a one-dimensional space, i.e., they can be collectively treated as one variable. As shown in Figure 5.9, with the two “interfering” variables  $b_1, b_2$  aligned into one variable  $b_1 + b_2$  it is possible to solve the three equations for the three effective variables and thereby recover the desired variables  $a_1, a_2$ .

The exact repair problem was solved through elegant and explicit alignment solutions for all  $(k, n)$  MDS codes such that  $n \geq 2k$  in [130, 137]. It was also shown in [130] that in general scalar linear codes do not meet the min-cut bound for the exact repair problem. The exact repair solution for  $(k, n)$  MDS codes in the low redundancy setting ( $k/n > 1/2$ ) was open, until recent work by Cadambe et al. in [29] and Suh and Ramchandran in [138] where the [CJ08] scheme was applied to obtain a complete solution. The general solution is summarized next with the help of Figure 5.10.

As shown in Figure 5.10 a file of size  $\Delta \times k$  symbols is to be stored in  $n$  distributed data storage nodes, each capable of storing  $\Delta$  symbols, such that the failure of any  $n - k$  nodes will not lead to permanent loss of information. Since only (any)  $k$  nodes need to survive, the total amount of data that must survive is no more than the size of the original file. We need the minimum storage property, i.e., we need MDS codes. The specific code constructed in Figure 5.10 consists of, first,  $k$  systematic nodes that simply store the  $k$  independent pieces of the file of size  $\Delta$  each without any coding, and then  $n - k$  parity nodes that

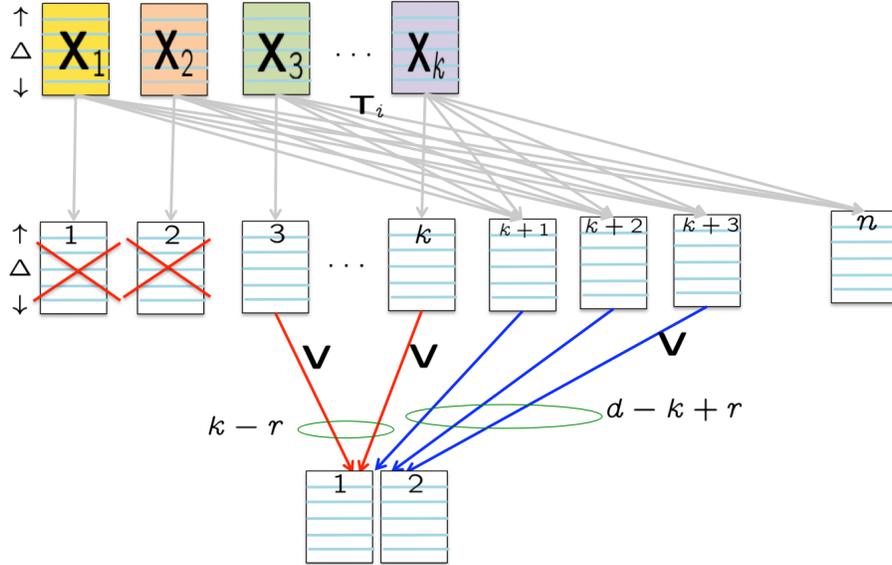


Fig. 5.10 [CJ08] scheme as used in [29] for the exact repair problem for an  $(k, n)$  MDS code when  $r = 2$  storage nodes fail and  $d$  of the remaining nodes are contacted by the repair agent.

store linear combinations. In particular, if we consider the file fragments  $X_1, X_2, \dots, X_k$  as  $k$  vectors each  $\Delta \times 1$ , i.e., each containing  $\Delta$  symbols, then each parity node contains a linear combination of all the fragments  $X_k$  each scaled by a *generic* diagonal matrix  $T_i$ . The MDS property of this code construction is established in [29]. The coding matrices  $T_i$  are analogous to the channel matrices in the wireless setting, and their diagonal nature is crucial for the application of the [CJ08] scheme.

Now suppose that  $r$  of the systematic nodes, e.g., nodes  $1, 2, \dots, r$  fail, where  $r \leq k$ , and a repair agent, capable of contacting  $d$  of the surviving  $n - r$  nodes, wishes to replicate the failed nodes. Note that we must have  $d \geq k$  to have any guarantee that the lost data can be recovered. Then what is the minimum number of linear equations that this repair agent must download from each of the  $d$  surviving nodes that it is able to contact, in order to replicate all the failed nodes? In particular, each node contains  $\Delta$  linear combinations of the file's data symbols, and the repair agent wishes to download  $\beta$  linear combinations, which may themselves be constructed as linear combinations

of the stored data available at the node being contacted. Thus, the repair agent needs to find the  $\beta \times \Delta$  matrix representing the projection of the  $\Delta$  dimensional object into  $\beta$  dimensions that it needs to download. This is the matrix  $\mathbf{V}^T$  and the setting is complete for [CJ08] scheme to be applied. Consistent with the [CJ08] scheme, the repair agent demands the same projection  $\mathbf{V}$  of the data stored in each of the  $d$  nodes that it is able to contact. Without loss of generality, we will assume that the  $d$  nodes contacted include among them the surviving  $k - r$  systematic nodes — this assumption can be justified by a simple change of basis operation.

To clarify the notation with an example, let us identify the corresponding terms in the simple example of Figure 5.9. For that example,  $k = 2, n = 4, \Delta = 2$ ,  $X_1 = (a_1, a_2)$ ,  $X_2 = (b_1, b_2)$  is the data stored in systematic nodes.  $X_3 = (a_1 + b_1, 2a_2 + b_2)$  and  $X_4 = (2a_1 + b_1, a_2 + b_2)$  can also be seen as

$$X_3 = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix} X_1 + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} X_2 \quad (5.1)$$

$$X_4 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} X_1 + \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix} X_2, \quad (5.2)$$

which highlights the *diagonal* nature of the combining matrices  $T_i$  used for the parity nodes. In this example,  $r = 1$  as node 1 fails,  $d = 3$  as nodes 2, 3, 4 are contacted. The matrix  $\mathbf{V}^T = [1 \ 1]$ , so that  $\beta = 1$ . The equation downloaded from the systematic node (Node 2) is simply  $\mathbf{V}^T X_2$  and the equations downloaded from the parity nodes (Nodes 3 and 4) are  $\mathbf{V}^T X_3, \mathbf{V}^T X_4$  respectively. The repair node only wants  $X_1$ , which means the information contained in  $X_2$ , i.e., symbols  $b_1, b_2$  only serves as interference. Finally, note that the equation downloaded from Node 2 serves as the perfect antidote for the interference terms  $(b_1, b_2)$  contained in the equations downloaded from the parity nodes because of interference alignment, i.e.,  $b_1, b_2$  always appear in the same linear combination  $b_1 + b_2$ .

Now let us return to the general setting of Figure 5.10. The original file data consists of  $k$  independent data vectors, each of size  $\Delta \times 1$ . Since the repair agent is only interested in  $r$  of these that were lost due to the failed systematic nodes, the remaining  $k - r$  systematic vectors

constitute undesired information for the repair agent, i.e., any projection of the information from these vectors constitutes interference. The desired information can then only come from the  $d - k + r$  parity nodes that are also contacted. These parity nodes contain linear combinations of all data vectors, including the interference vectors. Since the amount of data to be repaired is  $r \times \Delta$ , and this information comes only from the  $d - k + r$  parity nodes that are contacted, the number of equations to be downloaded from each parity node must be at least  $\frac{r\Delta}{d-k+r}$ . Suppose we download only this many equations, i.e.,  $\mathbf{V}$  has  $\beta = \frac{r\Delta}{d-k+r}$  columns representing the  $\beta$  projections to be downloaded from each contacted node. Note that in addition to the desired information we also download the interference that is mixed with it, and has the form  $\mathbf{V}^T \sum_{i=r+1}^k T_{\bullet} X_i$ . Since only its diagonal form is important and we wish to minimize notational clutter, we do not identify the specific subscript for the  $T$  matrices. In order to remove this interference we need “antidotes” from the  $k - r$  systematic nodes. The antidotes themselves take the form  $\mathbf{V}^T X_i$  for  $i = r + 1, \dots, k$ .

In order for the antidotes to be effective, we need them to match the interference. This leads to the requirement that:

$$\mathbf{V}^T \approx \mathbf{V}^T T_1 \approx \mathbf{V}^T T_2 \approx \dots \approx \mathbf{V}^T T_N. \quad (5.3)$$

This is the familiar problem that the [CJ08] scheme is designed to solve. Thus, for  $\Delta$  large enough, the [CJ08] scheme is able to align the antidotes almost perfectly with the interference so that all interference can be cancelled and the desired information recovered. The total data downloaded is  $d \times \beta$  and the data to be repaired is  $r\Delta$ . Thus, the normalized repair bandwidth with the scheme described above is computed as

$$\frac{\text{Data downloaded}}{\text{Data repaired}} = \frac{d\beta}{r\Delta} \quad (5.4)$$

$$= \frac{d}{d - k + r}, \quad (5.5)$$

where we made the substitution  $\beta = \frac{r\Delta}{d-k+r}$ . It turns out that this normalized repair bandwidth coincides with the information theoretic min-cut bound, thus establishing the optimality of the [CJ08] scheme as an asymptotic vector coding solution to the exact repair problem.

A few finer details are worth mentioning about the solution presented above. First, we chose the coding matrices  $T_i$  as “generic” matrices. Since we are dealing with finite fields, this is meant only to indicate a random choice of coding matrices over the finite field of interest. Further, the “random” choice of coding matrices is akin to Shannon’s random coding argument, ultimately intended not as a randomized construction with some probability of error, but as an argument leading to the proof of existence of good channel matrices over sufficiently large finite fields via the Schwartz–Zippel Lemma. Second, the arguments relying on field sizes being large enough carry the caveat that the required field size may be larger than the minimum requirement for MDS property. There is much room for improvement over the [CJ08] scheme in terms of field size requirements as well as nonasymptotic solutions, as explored most recently in [21, 113, 140]. Lastly, with regard to (5.3), note that the [CJ08] scheme only establishes an approximate equivalence between the antidotes and the actual interference. Therefore, some additional “equations” need to be downloaded as antidotes. However, this additional repair bandwidth expansion is a negligible fraction of the overall data downloaded, and is therefore inconsequential in the limit of interest (large file sizes).

*Open Problem:* Open problems in distributed storage that are related to interference alignment include finding practical codes to achieve the theoretical bounds achieved by the asymptotic [CJ08] scheme, finding optimal interference alignment solutions that trade-off storage capacity of the nodes (i.e., relaxing the minimum-storage constraint) against the repair bandwidth, and using interference alignment to efficiently detect and isolate corrupted storage nodes.

## 5.8 Multihop Interference Networks

Following the review of network coding applications, we return to wireless networks — specifically to multihop multiframe wireless networks. The reader may have noticed that all applications of interference alignment in wireless networks have focused on single hop networks. Indeed, applications of interference alignment to multihop interference

networks are not well understood and may be one of the exciting frontiers of research on interference alignment in the near future.

It is shown in [28] that relays do not increase the DoF of interference networks when the network is fully connected, i.e., all channel coefficients are nonzero and global channel knowledge is available, for almost all channel coefficient values. However, the use of relay nodes is known to simplify the achievable scheme, e.g., by enabling linear alignment schemes when lattice alignments would otherwise be needed and by requiring smaller symbol extensions when asymptotic schemes would otherwise be necessary [33, 48, 107, 110, 111].

Since relays are often used in settings where the direct signal is not strong enough, there is interest in layered multihop interference networks, especially layered two hop networks, where the sources can be heard by the relays and the relays are heard by the destinations but the sources cannot be heard directly by the destination nodes. The discussion in this section is primarily restricted to layered interference networks.

DoF of two hop layered interference networks are known when there is a relay that is equipped with sufficiently many antennas [16, 141] or when there is a large number (in many cases, approaching infinity) of distributed relays [11, 15, 38, 98, 122]. In all cases the goal is to completely eliminate interference between communicating transmitter–receiver pairs. With a MIMO relay this can be accomplished by viewing the network as a MAC channel from the distributed sources to multiple-antenna relay followed by a BC channel from the multiple antenna relay to the distributed destinations. With distributed relays, the zero forcing of interference must be accomplished in a distributed fashion. The distributed zero forcing of interference is known by various names in literature including *distributed orthogonalization*, *multi-user zero-forcing* and *interference neutralization*, all of which refer to the same essential idea. An excellent overview of these techniques is presented in [11].

Of particular interest is the setting with a *finite* number of *distributed* relay nodes, especially when the number of relay nodes is not enough to allow perfect distributed orthogonalization. For the  $K \times R \times K$  setting, i.e., a  $K$  user layered interference channel where communication takes place through  $R$  distributed relay nodes, Rankov

and Wittneben have shown in [122] that a necessary condition for distributed orthogonalization is that  $R \geq K(K - 1) + 1$ , which is identified as the *minimum relay configuration*. For example, in a 2 user interference channel with 3 relay nodes, it is possible to eliminate all interference. The overconstrained network with fewer relays, the  $2 \times 2 \times 2$  interference channel, therefore becomes particularly interesting. On the one hand the network min-cut outer bound is still 2 DoF, but the number of relays is evidently smaller than the minimum relay configuration identified in [122] for distributed orthogonalization. Here we review two distinct DoF-optimal approaches for this channel, obtained in [51, 71]. While the first approach is opportunistic, the second approach shows that alignment of interference plays an interesting role in facilitating the ultimate neutralization of interference without requiring opportunistic schemes.

Figure 5.11 shows the  $2 \times 2 \times 2$  interference channel, where two transmitters wish to send independent messages to their corresponding receivers over a concatenation of two interference channels. While several studies of this channel [135, 32] treat the two hops as interference channels, it is evident that any such approach cannot achieve more than 1 DoF (because the 2 user interference channel has only 1 DoF). On the other hand, by treating it as a concatenation of two X channels, Cadambe and Jafar have shown in [23] that  $4/3$  DoF can be achieved almost surely. This requires that the sources split their messages into two independent parts, one for each relay. The four independent messages transform the  $2 \times 2 \times 2$  IC into the X channel setting over each hop and  $4/3$  DoF become achievable.

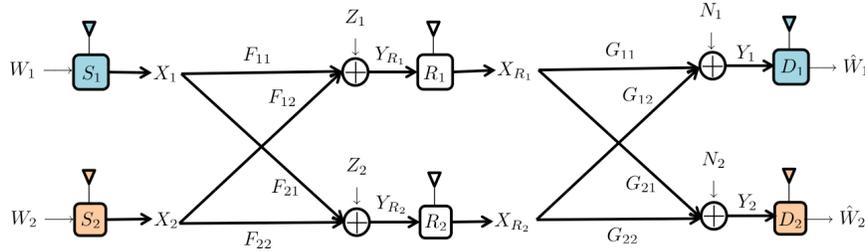


Fig. 5.11 The  $2 \times 2 \times 2$  interference channel.

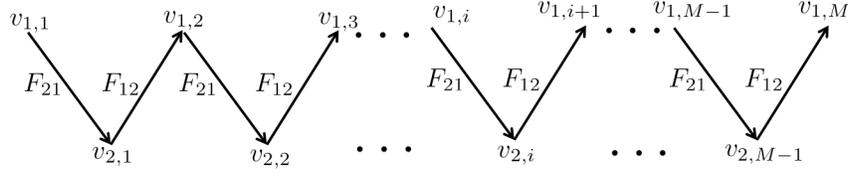
An interesting idea called *opportunistic interference neutralization* is introduced by Jeon et al. in [71]. *Opportunistic* interference neutralization refers to the idea that the relays buffer the symbols from the first hop and transmit them over the second hop only when the channel matrices  $F$  and  $G$  for the two hops are complementary in the sense that  $GF$  is a diagonal matrix. Essentially, the matrices invert each other so that all interference, which corresponds to cross-terms in the product matrix, is eliminated. For a fairly broad class of channels, including, e.g., i.i.d. Rayleigh fading channels, Jeon et al. show in [71] that complementary matrices can be matched almost perfectly so that the DoF = 2 can be achieved for the  $2 \times 2 \times 2$  IC. It is remarkable that the network DoF min-cut outer bound of two is achievable for such channels, considering that for one hop interference channels, no more than half the min-cut is achievable [22, 57]. While the idea of opportunistic interference neutralization is very powerful, like ergodic interference alignment, it is limited to a special class of distributions and in particular does not apply to constant channels.

For constant channels, Gou et al. [51] have recently introduced the idea of *aligned interference neutralization* that is shown to achieve the DoF = 2 outer bound for the  $2 \times 2 \times 2$  IC for almost all channel values. We summarize the aligned interference neutralization scheme below. For simplicity of exposition, and without loss of generality, let us normalize  $F_{11} = F_{22} = G_{11} = G_{22} = 1$ .

Signaling is performed over  $M$  rational dimensions, within which  $M$  signals from User 1 and  $M - 1$  signals from User 2 are sent, so that a total of  $\frac{2M-1}{M}$  DoF are achieved. Since  $M$  can be made arbitrarily large, the DoF min-cut outer bound is achieved in the limit of large  $M$ . Up to a scaling factor needed to satisfy power constraints, the transmitted signals are constructed as

$$X_1 = \sum_{m=1}^M v_{1,m} x_{1,m} \quad (5.6)$$

$$X_2 = \sum_{m=1}^{M-1} v_{2,m} x_{2,m} \quad (5.7)$$


 Fig. 5.12 The  $2 \times 2 \times 2$  IC.

so that  $v_{i,m}$  are the rational “beamforming” directions and  $x_{i,m}$  are the  $Q$ -ary coded symbols as explained in Section 4.7. The  $v_{i,m}$  directions are chosen according to Figure 5.12, so that  $v_{2,i} = F_{2,1}v_{1,i}$  and  $v_{1,i+1} = F_{1,2}v_{2,i}$ .

Any choice of  $v_{1,1}$  therefore determines all  $v_{k,i}$ . The alignment of signals along rational dimensions at Relays 1 and 2 is illustrated in the table below:

At Relay 1:

dimension	$v_{1,1}$	$v_{1,2}$	$v_{1,3}$	$\cdots$	$v_{1,i+1}$	$\cdots$	$v_{1,M}$
dimension		$F_{12}v_{2,1}$	$F_{12}v_{2,2}$	$\cdots$	$F_{12}v_{2,i}$	$\cdots$	$F_{12}v_{2,M-1}$
symbol	$x_{1,1}$	$x_{1,2} + x_{2,1}$	$x_{1,3} + x_{2,2}$		$x_{1,i+1} + x_{2,i}$		$x_{1,M} + x_{2,M-1}$

At Relay 2:

dimension	$F_{21}v_{1,1}$	$F_{21}v_{1,2}$	$\cdots$	$F_{21}v_{1,i}$	$\cdots$	$F_{21}v_{1,M-1}$	$F_{21}v_{1,M}$
dimension	$v_{2,1}$	$v_{2,2}$	$\cdots$	$v_{2,i}$	$\cdots$	$v_{2,M-1}$	
symbol	$x_{1,1} + x_{2,1}$	$x_{1,2} + x_{2,2}$		$x_{1,i} + x_{2,i}$		$x_{1,M-1} + x_{2,M-1}$	$x_{1,M}$

Each relay makes hard decisions on the *aligned*  $Q$ -ary symbols. So, Relay 1 demodulates  $x_{1,1}, x_{1,2} + x_{2,1}, x_{1,3} + x_{2,2}, \dots, x_{1,i+1} + x_{2,i}, \dots, x_{1,M} + x_{2,M-1}$  while Relay 2 demodulates symbols  $x_{1,1} + x_{2,1}, x_{1,2} + x_{2,2}, \dots, x_{1,i} + x_{2,i}, x_{1,M-1} + x_{2,M-1}, x_{1,M}$ .

These demodulated symbols are sent over the second hop from each relay, again in an aligned fashion similar to the first hop but with phase reversals chosen to cancel undesired symbols, as shown in Figure 5.13 where the top row of symbols corresponds to Destination 1 and the bottom row corresponds to Destination 2. Thus, for example, symbol  $-x_{1,1}$  aligns with symbol  $x_{1,1} + x_{2,1}$  at Destination 2 leaving just the desired symbol  $x_{2,1}$ , symbol  $-(x_{1,i} + x_{2,i-1})$  aligns with  $x_{1,i} + x_{2,i}$  leaving just the desired symbols  $x_{2,i} - x_{2,i-1}$ . The chain of desired symbols can be successfully demodulated once all interfering symbols are removed. At the same time, the alignments at Destination 1 force the symbol  $-(x_{1,1} + x_{2,1})$  to align with  $x_{1,2} + x_{2,1}$ , leaving only

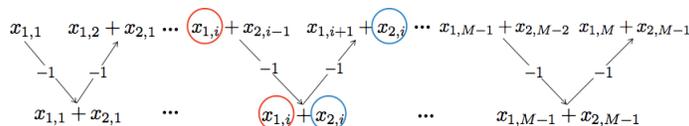


Fig. 5.13 Aligned interference neutralization over the second Hop.

the desired symbols  $x_{1,2} - x_{1,1}$ , the symbol  $-(x_{1,i} + x_{2,i})$  aligns with  $x_{1,i+1} + x_{2,i}$  leaving only the desired symbols  $x_{1,i+1} - x_{1,i}$  and so on, so that the chain of desired symbols can be successively demodulated. The error probability of these demodulations approaches zero for high SNR according to the rational dimensions framework, so that the *coded* rate for which probability of error can be made arbitrarily small approaches a pre-log factor of  $1/M$  per rational dimension. Thus, a total DoF of  $\frac{2M-1}{M}$  are achieved, and in the limit of large  $M$ , the network min-cut outer bound value of DoF = 2 is achieved as well. Gou et al. also show that when restricted to linear beamforming techniques and constant complex channel values, asymmetric complex signaling provides the  $M = 2$  vector dimensions along which aligned interference neutralization can be used to achieve  $\frac{2M-1}{M} = 1.5$  DoF for almost all channel coefficient values, which is remarkably higher than the best known previous result of  $4/3$  DoF possible by viewing the two hops as  $X$  channels. The similarity of the alignment scheme with the [CJ08] scheme is also remarkable, as they are both asymptotic and construct signaling dimensions out of the products of integer powers of the channel coefficients.

Shomorony and Avestimehr in [134] initiate the DoF study of multihop layered interference networks with arbitrary connectivity within each hop. With two source nodes  $S_1, S_2$ , two destination nodes  $D_1, D_2$ , arbitrary number of relay nodes at each layer and arbitrary connectivity between each pair of adjacent layers, there are four possible independent unicast messages  $W_{ij}, i, j \in \{1, 2\}$ , such that  $W_{ij}$  is the message originating at source  $S_i$  and intended for destination  $D_j$ . If the message set is restricted to  $W_{11}, W_{22}$ , we have a two-user multihop layered interference channel. For this channel Shomorony and Avestimehr show that for almost all channel coefficient values corresponding to all possible layered network graphs, the DoF values can only belong to

the set  $\{1, \frac{3}{2}, 2\}$ . An important concept of *manageable interference* is introduced, which makes precise the condition under which interference arriving from multiple paths can be neutralized. Also intriguing are the wireless counterparts of the *butterfly* and *grail* topologies previously highlighted for wired multiple unicast setting [30, 132, 158]. However, interference alignment appears to be essential only for the  $2 \times 2 \times 2$  setting described above, where the aligned interference neutralization scheme of Gou et al. [51] is utilized over the rational dimensions framework of [99].

Wang et al. in [154] explore the 2-source 2-sink layered multihop network DoF when all possible unicast messages  $W_{ij}, i, j \in \{1, 2\}$  are allowed (also known as the X network setting). In this case, with arbitrary number of relays at each hop, and arbitrary connectivity between adjacent layers, the DoF are shown to belong to the set  $\{1, 4/3, 3/2, 5/3, 2\}$ . Interference alignment plays a much more significant role in this setting. Figure 5.14 shows two networks with DoF values  $3/2$  and  $5/3$  respectively, where interference alignment is used to establish the achievability. The network on the left in Figure 5.14 has  $3/2$  DoF in the X setting but can only achieve 1 DoF in the interference channel setting where only two messages (either  $W_{11}, W_{22}$ , or  $W_{12}, W_{21}$ ) are allowed. The network on the right in Figure 5.14 has  $5/3$  DoF in the X setting but has only 1 DoF as the interference channel with messages  $W_{11}, W_{22}$ , and only  $3/2$  DoF as the interference channel with messages  $W_{12}, W_{21}$ . The interference alignment schemes for the two networks in the X setting are illustrated in the figure.

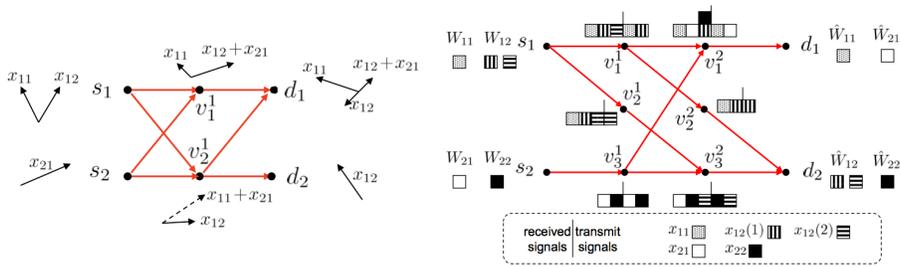


Fig. 5.14 Interference alignment in 2-source 2-sink layered X networks.

*Open Problem:* While the aligned interference neutralization scheme is shown to achieve the network min-cut DoF outer bound of the  $2 \times 2 \times 2$  IC, the DoF for extensions such as  $3 \times 3 \times 3$  IC are not known except in the opportunistic interference neutralization setting. Characterizations such as GDoF, capacity within constant gap etc. are unknown for almost all multihop interference networks, with few exceptions [97]. These are promising research avenues for future studies on the benefits of interference alignment in wireless networks and may also provide useful insights for solving the multiple unicasts network coding problems which are also concerned with multihop interference management issues.

## 5.9 Bidirectional Relay Interference Networks

Continuing with the theme of multihop interference networks, let us now add the element of bidirectional communication. Two-way communication with the help of one or more relay nodes is a topic that has recently drawn much attention [4, 11, 28, 76, 77, 87, 121, 122, 129]. Remarkably, signal alignment is a key element of achievable rate strategies of bi-directional relay interference networks. The alignment of signals in this setting is explored in the signal level dimension with the use of deterministic channel models and single antenna nodes by Sezgin et al. in [129] and in the signal vector space dimension with the use of linear beamforming with multiple antenna nodes by Lee et al. in [87]. Here we summarize the signal vector space alignment approach of Lee et al. where the need for alignment is evident from a DoF perspective. By analogy to the two-user case of Analog Network Coding [76], we label this problem the Aligned Analog Network Coding problem in Figure 5.15.

The setting illustrated in Figure 5.15 consists of three nodes, each equipped with two antennas, and six independent messages  $W_{ij}, i, j \in \{1, 2, 3\}$  such that message  $W_{ij}$  originates at node  $j$  and is intended for node  $i$ . There are no direct links between these nodes. All communication must pass through a relay node that is equipped with three antennas. The channel coefficients are generic and may be regarded as fixed. Communication takes place in two phases — an uplink phase

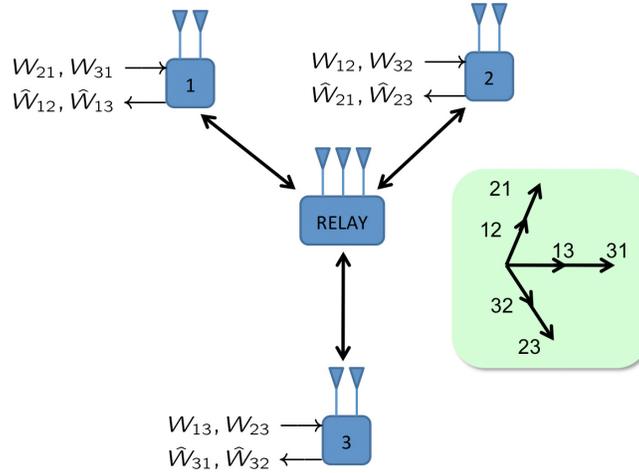


Fig. 5.15 Aligned analog network coding [87].

where all nodes transmit to the relay node and a downlink phase where the relay transmits to all the nodes.

It is shown by Lee et al. in [87] that each of the six messages can achieve 1 DoF over one channel use of the uplink phase and one channel use of the downlink phase. The need for alignment is understood as follows. Consider node 1 and the downlink phase. There are six messages mixed into the three-dimensional space available to the relay. Out of these six messages there are two messages —  $W_{21}, W_{31}$  — that are harmless to Node 1 in the sense that these messages originated at Node 1, so they are known to Node 1, and therefore any interference from these messages can be cancelled. That leaves four messages out of which the two messages  $W_{12}, W_{13}$  are desired by Node 1, so that these messages must not align, i.e., they must occupy one dimension each at the relay. This leaves one dimension at the relay node and two potentially harmful interfering messages —  $W_{23}, W_{32}$ . Therefore  $W_{23}$  and  $W_{32}$  must align into one dimension. Arguing similarly for each node, it is evident that  $W_{ij}$  and  $W_{ji}$  should align into one dimension at the relay. Thus, six messages align into the three dimensions available to the relay node. There remain two questions. First, is this alignment possible?, and second, do the alignments allow each node to recover its desired message?

To answer the first question, consider, e.g., messages  $W_{23}, W_{32}$ , originating at nodes 2,3 that need to align at the relay node. Nodes 2 and 3 each have two antennas while the relay has three antennas. In other words, nodes 2 and 3 can each access a two-dimensional vector subspace out of the three-dimensional vector space visible to the relay node. Under generic channel conditions, these two subspaces must intersect in  $2 + 2 - 3$  dimensions, i.e., there is a one-dimensional subspace that must be accessible from both Node 2 and from Node 3. It is this one-dimensional vector subspace, i.e., vector along which the signals for messages  $W_{23}$  and  $W_{32}$  will align. In fact this is the only vector subspace along which such alignment is possible. Similarly, nodes 1, 2 cast two-dimensional shadows within the three-dimensional subspace seen by the relay. These shadows must overlap in one dimension, which will be the direction chosen by Nodes 1 and 2 to align their respective messages  $W_{12}, W_{21}$ , respectively. The same argument applies to the alignment of messages  $W_{13}, W_{31}$  as well.

To answer the second question, we elaborate upon the transmission scheme used by the relay in the downlink phase. The relay separates the signal dimension along which  $W_{23}, W_{32}$  are aligned and rotates it into a vector orthogonal to the two-dimensional subspace corresponding to the  $2 \times 3$  downlink channel of Node 1. Similarly, the relay separates the signal dimension along which  $W_{12}, W_{21}$  are aligned and rotates it into the direction of the zero forcing vector for Node 3 and separates the signal dimension along which  $W_{13}, W_{31}$  are aligned and sends these aligned messages along the zero forcing vector of Node 2. Note that since the relay has three antennas and each node has only two antennas, there is a unique zero-forcing direction for each node available to the relay, i.e., the one-dimensional null space of the  $2 \times 3$  channel matrix for that node. By using these zero-forcing directions for the undesired aligned messages, the relay ensures that in the downlink phase the desired messages are resolvable as follows. Due to generic channel conditions, Node 1 sees two linearly independent signal vectors, one containing aligned messages  $W_{12}, W_{21}$  and the other containing aligned messages  $W_{13}, W_{31}$ . Since Node 1 already knows messages  $W_{21}, W_{31}$  and all the channel coefficients it is able to remove the contribution from these messages. Thus, Node 1 sees two linearly independent signals, one

containing the message  $W_{12}$  and the other containing the message  $W_{13}$ . These are the two desired messages for Node 1, and thus, Node 1 is able to resolve the two desired messages within the two-dimensional signal space available to it. By symmetry of construction, the same argument is applied to each node. In this manner each of the six messages is able to reach its desired destination along a linearly independent dimension over one channel use of the uplink and one channel use of the downlink, achieving a total of 6 DoF over these two channel uses.

*Open Problem:* There is an abundance of open problems in the bidirectional interference networks setting. From a DoF perspective, the open problems include the DoF with generalized antenna configurations in the setting of Figure 5.15 and the DoF with distributed relay nodes. GDoF characterizations for MIMO bidirectional relay networks are particularly interesting because of the need to combine the signal level alignments shown to be necessary in [129] and the signal space alignments shown to be necessary in [87].

## 5.10 Cooperative Interference Networks

Cooperation among nodes is widely recognized as a way to alleviate interference in wireless networks. Instead of attempting a review of the vast body of literature on cooperative communications we will point the reader to excellent survey and tutorials such as [43, 84]. Our goal in this section is to highlight some interesting aspects that primarily concern interference alignment in cooperative interference networks.

If all channels are of comparable strength, then it is shown in [28, 57] that there is no DoF benefit from cooperation. Thus, the benefits of cooperation, at the DoF level, are enabled by the existence of heterogeneous links, such as backhaul links that do not compete for wireless spectrum or stronger links between local clusters of nodes over a certain hop that may be taken as ideal links if the bottleneck for the end to end communication rate lies elsewhere. These ideal links may enable cognitive settings, i.e., where some messages can be shared between nodes. The presence of cognitive and cooperative transmitters and receivers in a wireless interference network along with the localized nature of wireless connectivity creates many interesting problems where not only does

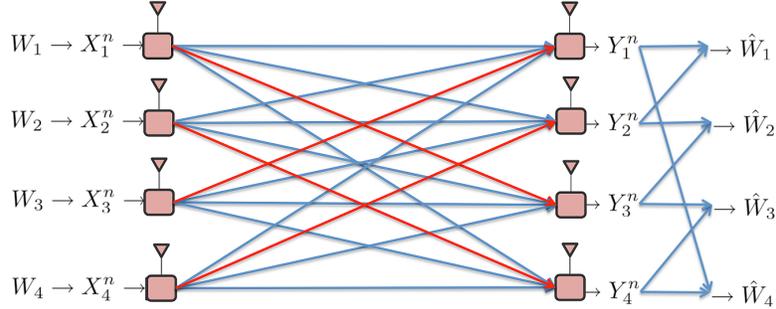


Fig. 5.16 Pairwise cooperative decoding under different connectivity models.

interference alignment play a significant role, but also fundamentally new aspects of interference alignment may come to light.

An example of a cooperative interference network is illustrated in Figure 5.16 where four transmitters send four independent messages  $W_1, W_2, W_3, W_4$  to four receivers. In decoding the messages, cooperation is allowed between adjacent receivers in the sense that the decoder for message  $W_1$  has access to the received signals  $Y_1^n, Y_2^n$  from receivers 1 and 2, respectively; the decoder for message  $W_2$  has access to  $Y_2^n, Y_3^n$ ; the decoder for message  $W_3$  has access to  $Y_3^n, Y_4^n$ ; and the decoder for message  $W_4$  has access to  $Y_4^n, Y_1^n$ . In this network, if all channel coefficients are nonzero, then Annapureddy et al. in [2] find the DoF to be  $2/3$  per user, similar to the SIMO interference channel where each receiver has two antennas. The achievability scheme is based on the [CJ08] alignment scheme, but involves several challenging aspects because of the shared antennas which create spatial dependencies between the channel coefficients of the equivalent SIMO interference channel. In other words, the independence between desired channel coefficients and interference carrying channel coefficients, the basis for the assumption that desired signals do not overlap with interference, does not hold when receive antennas are shared by decoders. In fact if this setting is naturally extended to more than four users then it can be shown that there is a loss of DoF relative to the corresponding SIMO interference channel. While the SIMO interference channel with two antennas at each receiver continues to achieve  $2/3$  DoF per user, the DoF of the fully connected five user interference channel with pairwise cooperative

decoders is bounded above by  $3/5$  DoF per user. Another surprising aspect of this setting, explored by Wang et al. in [155] is the role of channel connectivity. Specifically, suppose in Figure 5.16 the channel coefficients for all the red links are replaced by zeros, i.e., these links are disconnected. Note that these links do not carry any desired message information. For example, the red link from Transmitter 4 to Receiver 2 carries only message  $W_4$  which is not desired by the decoder for either message  $W_1$  or  $W_2$ , i.e., the two decoders that have access to Receiver 2. Thus, one might expect that setting this link to zero will not reduce the DoF. But as shown by [155] setting the red links to zero forces a DoF outer bound of  $\frac{3}{5}$  per user, which is strictly smaller than the  $2/3$  DoF per user achievable in the fully connected setting of Figure 5.16. Interesting insights such as this, as well as the novel alignment schemes as well as new outer bounds needed to account for shared antennas, make this problem of great fundamental theoretical interest. The reciprocal setting of cooperative transmitters is similarly interesting.

Another application of interference alignment in wireless networks with wired backhaul links is the idea of *interference alignment and cancellation* (IAC) introduced in the work of Gollakota et al. in [47]. An example of the IAC framework is illustrated in Figure 5.17, where a two user MIMO interference channel is shown with all nodes equipped with three antennas. In the absence of receiver cooperation it is known that this channel can achieve only 3 DoF [65]. Now suppose there is a mechanism through which the receivers can share their *decoded messages*, i.e., once a message is decoded at a receiver the message bits are communicated to the other receiver who can reconstruct the

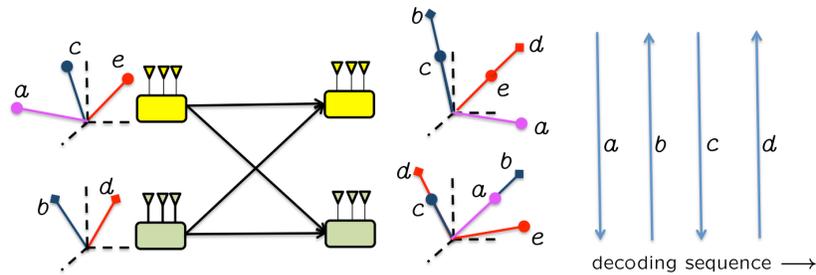


Fig. 5.17 Interference alignment and cancellation [47].

interference due to that message and cancel it from its own received signal. Under this assumption, Figure 5.17 shows how User 1 can achieve 3 DoF while User 2 simultaneously achieves 2 DoF through IAC. Specifically, User 1 communicates three messages, labeled as  $a, c, e$  in Figure 5.17, while User 2 communicates two messages, labeled  $b, d$ . The signaling dimensions for these messages are picked so that  $a, b$  align along the same signal vector at Receiver 2;  $b, c$  align along one signal vector at Receiver 1;  $c, d$  align along one vector at Receiver 2; and  $d, e$  align along one vector at Receiver 1.  $a$  does not align with anything at Receiver 1. As shown in Figure 5.17 this creates three signaling vectors at each receiver, which can be resolved from each other, but each vector may contain up to two aligned messages. The decoding proceeds in the sequence  $a, b, c, d, e$ . First, Receiver 1 decodes message  $a$ , which can be resolved because it is not aligned with anything else. This decoded message is made available to Receiver 2 through a backhaul link, allowing Receiver 2 to cancel interference from  $a$ . Cancelling  $a$  allows Receiver 2 to decode message  $b$ , which was aligned only with  $a$ . The decoded message  $b$  is then passed back to Receiver 1 who cancels the interference from  $b$  to obtain an uninterfered dimension for its desired message  $c$ . After Receiver 1 decodes message  $c$  this message is made known to Receiver 2 who cancels interference from this message to obtain a clean signal vector for decoding message  $d$ . Finally, message  $d$  is provided to Receiver 1 who cancels the interference due to it and obtains a clean signal dimension for message  $e$ , thus concluding the decoding sequence.

Multiple antennas are not essential for the IAC scheme described above. For example, the scheme can be similarly operated for a single antenna two user interference channel with time-varying/frequency-selective channel coefficients. Over an  $M$  dimensional signal space obtained by considering  $M$  time-slots or  $M$  frequency slots, User 1 can obtain  $M$  DoF and User 2 can simultaneously obtain  $M - 1$  DoF with the chain of alignments and cancellations described above, so that a total of  $\frac{2M-1}{M}$  DoF are achieved. Note that the achieved DoF approach the value 2, which is the maximum DoF possible with full cooperation, i.e., joint signal processing between the two receivers. Also remarkable are the parallels between the chain of alignments and cancellations

described above, and the chain of alignments and cancellations used in the aligned interference neutralization scheme of [51] presented earlier.

*Open Problem:* An interesting open problem is to characterize the DoF of the  $K$  User interference channel with clustered cooperative decoders, such as the setting illustrated in Figure 5.16 with  $K > 4$  users. The DoF with Wyner connectivity in this setting are also interesting, as mentioned earlier, due to the new fundamental issues that arise with shared antennas. The DoF duality of transmitter cooperation and receiver cooperation is also an open problem.

### 5.11 Secrecy

As modern communication systems rely increasingly on wireless transmissions, security of the communicated information is a growing concern. Fortunately, the idea of interference alignment leads to naturally secure communication schemes. This is pointed out first by Koyluoglu et al. in [83] where the secure DoF of the  $K$  user interference channel are determined. Combining random binning arguments needed for secrecy along with the [CJ08] scheme needed for alignment it is shown that the  $K$  user interference channel can achieve  $\frac{K-2}{2K-2}$  secure DoF per user. The intuition for robustness to secrecy demands can be understood as follows. As we know, the goal of interference alignment is to allow the resolution of desired symbols even when the number of symbols is much larger than the number of equations, by aligning the undesired symbols into the smallest possible subspace. While the desired symbols are resolvable, interference alignment forces the signal dimensions of undesired symbols to overlap, essentially forcing undesired transmissions to jam each other, thereby naturally ensuring that the undesired information is not decodeable. Similarly, with an external eavesdropper whose channel state may not be known to the transmitters, Koyluoglu et al. show in [83] that secure DoF of  $\frac{K-2}{2K}$  are achievable per user. The intuition here may be understood as follows. Since no effort is made to align signals for the eavesdropper, the eavesdropper sees essentially generic linear equations in the information carrying symbols. Since the number of symbols far exceeds the number of equations, none of the symbols can be resolved by the eavesdropper.

The probability of symbols accidentally aligning when passing through generic channels is zero, and therefore security against the eavesdropper is naturally available. In [49], Gou and Jafar develop corresponding secure DoF results for the X network setting with  $M$  transmitters,  $N$  receivers and  $MN$  independent messages, one from each transmitter to each receiver. Here also the key ingredients are the random binning scheme for secrecy in conjunction with the [CJ08] scheme for interference alignment.

While Koyluoglu et al. in [83], as well as Gou et al. in [49] focus on time-varying/frequency-selective channels, DoF benefits of interference alignment for the  $K$  user interference channel with *constant* channels under secrecy constraints and certain symmetric assumptions are shown by He and Yener in [56]. The key to He and Yener's work is the idea of layered lattice alignment, previously applied by Sridharan et al. in [136] to the same setting without secrecy constraints. A novel aspect of [56] is the use of nested lattice codes instead of the sphere-shaped lattice codes used in [136], which allows He and Yener to achieve higher DoF than possible with the approach of Sridharan et al.

The secure DoF of the  $K$  user Gaussian MAC in the presence of an eavesdropper, is explored by Bagherikaram et al. in [6]. It is shown that secure DoF of  $\frac{K-1}{K}$  is achievable, almost surely, by aligning the signals at the eavesdropper so that they are not resolvable, while keeping the signals distinguishable at the desired receiver. This can be achieved through Gaussian codebooks and linear beamforming schemes if the channels are time-varying/frequency-selective and through lattice alignment within the rational dimensions framework of [99] if the channel coefficient values are fixed.

The two user MAC with an external eavesdropper is considered by Bassily and Uluks in [8] under the assumption of time varying channels. Two schemes, called scaling based alignment (SBA) and ergodic secret alignment (ESA) are proposed, both of which are capable of achieving 1/2 secure DoF per user, and both of which exploit time-variations to achieve the alignment of signals at the eavesdropper. In particular, the ESA scheme is motivated by the ergodic interference alignment scheme of Nazer et al. [101] and achieves the stronger rate performance of the two schemes. As usual, these schemes have

the advantage of providing explicit DoF guarantees, unlike the implicit DoF characterizations of rational dimensions framework which can be shown to hold for almost all channel coefficients but cannot be explicitly applied to any given values of channel coefficients.

Interference alignment is also the key to the secure DoF characterization of a MISO compound wiretap channel explored by Khisti in [80]. In this setting there is a sender equipped with  $M$  antennas and two sets of  $J_1$  and  $J_2$  receivers, and each receiver is equipped with only one antenna. If one set of receivers is the legitimate set while the other consists of eavesdroppers and the channels are generic, then Khisti shows in [80] that secure DoF of 1 is achievable if  $\min(J_1, J_2) < M$  and secure DoF of  $\frac{M-1}{M}$  is achievable if  $\min(J_1, J_2) \geq M$ . Another variant of this problem is the compound private BC channel where each set of receivers wants a message that should be kept secret from the other set. In this setting, [80] shows that secure DoF of  $2, 2\frac{M-1}{M}, 2\frac{M-1}{M+1}$ , are possible when  $M$  is larger than  $\max(J_1, J_2)$ , between  $\max(J_1, J_2)$  and  $\min(J_1, J_2)$ , or smaller than  $\min(J_1, J_2)$ , respectively. The achievable scheme is based on the application of [CJ08] scheme in the rational dimensions framework of [99], and is similar to the compound MISO BC channel setting explored earlier in [52, 93] without secrecy constraints.

# 6

---

## Conclusion

---

Interference alignment is a relatively new idea in network information theory that has enabled significant progress in our understanding of the capacity limits of many wireless and wired communication networks. Correctly accounting for the number of signaling dimensions — the degrees of freedom — is one of the first steps toward a capacity characterization. Interference alignment has shown that even at this coarse level, theoretical limits can be orders of magnitudes higher than previously thought possible, especially for large networks.

While, due to their asymptotic high SNR nature, DoF (and GDoF) characterizations may seem of limited practical relevance, the surprising insights regarding new achievable schemes that emerge out of DoF studies are their most useful and far-reaching outcomes. Put simply, DoF characterizations may not provide us the recipe for a capacity achieving scheme, but they do serve to identify the main ingredients — often surprising ones. To that extent, interference alignment has been most successful as the connecting thread across a burst of innovative techniques that include new ways to exploit symbol extensions, channel variations, asymmetric complex signaling, structured codes, rational

dimensions, opportunistic transmissions, channel correlations, delayed feedback and reconfigurable antennas, as discussed in Sections 4 and 5.

The theoretical understanding of interference alignment is still not fully developed. It is limited primarily to linear channel models like Gaussian networks and linear deterministic frameworks. A representation of interference alignment in the classical discrete memoryless framework is not yet available, nor is it clear if such a representation should exist. Indeed, with increasing evidence on the side of structured coding techniques (in part due to the role of interference alignment), the existence of single-letter capacity characterizations is itself in question. Further, within the class of linear channel models there are extremely interesting fundamental questions that are not yet conclusively answered. For instance, is the bandwidth expansion aspect of interference alignment schemes as noted in [55] fundamentally inescapable? What significance does the rational/irrational nature of channel coefficients have at finite SNR? Are linear schemes always DoF optimal for generic time-varying channels and is the equivalence between this setting and the rational dimensions framework for constant channels an absolute equivalence? and are there new information theoretic inequalities hidden in the fundamental rules for simultaneous unions and intersections of multiple linear vector subspaces? These questions as well as the open problems mentioned throughout the paper are only a small sample of the many intellectually intriguing issues to be addressed by future work that are all centered around interference alignment.

While in this monograph we focused primarily on theoretical issues, there is an increasing body of work that explores practical challenges faced in the implementation of available theoretical interference alignment schemes. These challenges include the overhead for acquiring enough channel knowledge, the penalty of residual channel uncertainty at the transmitters, impact of channel correlations, tracking the interference alignment solution under time varying channels etc., and have been investigated by Thukral and Boelcskei in [144], Yu et al. in [164], Niu and Haimovich in [108], Krishnamachari and Varanasi in [85], El Ayach and Heath in [5], Peters and Heath in [117], Nosrat-Makouei et al. in [109], and Lee et al. in [88]. As new theoretical ideas related

to interference alignment continue to emerge, it is expected that they will similarly be put to test to measure their practical impact. It is this convergence of theoretical and practical interests that continues to accelerate the pace of progress on interference alignment schemes and is likely to produce many more breakthroughs in the years to come.

# A

---

## Degrees of Freedom (DoF)

---

In a communication network with  $m$  independent messages  $W_1, W_2, \dots, W_m$ , the rate tuple  $(R_1, R_2, \dots, R_m)$  is said to be achievable if there exists a sequence of codebooks (with increasing code lengths) such that the probability that any of the messages is decoded in error by its desired destination, can be made arbitrarily small by choosing code-words that are long enough. The closure of the set of all achievable rates is called the capacity region. In Gaussian networks, the capacity region depends on the local additive white Gaussian noise at each receiver, the transmitted signal power available to each transmitter, and the channel coefficient values that scale each signal as it propagates from the transmitters to the receivers. The DoF metric is primarily concerned with the limit where the total transmit power approaches infinity, while the values of channel coefficients and the local noise power remain unchanged. Thus, if we denote by  $C(P)$  the sum capacity with the total transmit power  $P$ , then the DoF metric  $\eta$  is defined as

$$\eta = \lim_{P \rightarrow \infty} \frac{C(P)}{\log(P)}, \quad (\text{A.1})$$

which is also equivalently stated as

$$C(P) = \eta \log(P) + o(\log(P)), \quad (\text{A.2})$$

where the  $o(\log(P))$  term is some function  $f(P)$  such that

$$\lim_{P \rightarrow \infty} \frac{f(P)}{\log(P)} = 0. \quad (\text{A.3})$$

It is worth pointing out as a cautionary note that it has not been established that the limit involved in the DoF metric *must* exist for all Gaussian networks with full CSI. Nevertheless, the significance of the DoF metric is widely recognized as a measure of the accessible signal dimensions in a communication network. Next, we elaborate further on this interpretation of DoF.

Consider a point-to-point Gaussian channel:

$$Y = HX + N,$$

where over a channel use  $Y$  is the output symbol,  $H$  is the channel coefficient,  $X$  is the channel input symbol, and  $N$  is the additive white gaussian noise (AWGN) term. All symbols are complex. The input is subject to the power constraint  $E[|X|^2] \leq P$  and the  $N$  is i.i.d circularly symmetric complex Gaussian  $\mathcal{N}^c(0, \sigma^2)$  across channel uses. The capacity of this AWGN channel was shown by Shannon to be

$$C = \log \left( 1 + P \frac{|H|^2}{\sigma^2} \right) \quad (\text{A.4})$$

bits per channel use (the log is assumed to be to the base 2). It follows then that

$$C = \log(P) + o(\log(P)). \quad (\text{A.5})$$

So this channel has one DoF. It is important to note that the channel strength  $H$  and the noise power  $\sigma^2$  are not relevant in this approximation, since they do not scale with  $P$ .

If we have  $M$  parallel AWGN channels

$$Y_m = H_m X_m + N_m$$

with power constraint

$$\frac{1}{M} \sum_{m=1}^M \mathbb{E}[|X|_m^2] \leq P \quad (\text{A.6})$$

and i.i.d. noise terms with noise powers  $\sigma_m^2$ , where  $m \in \{1, 2, \dots, M\}$ , and all channel coefficients are non zero then it is easy to see that the total capacity of these channels

$$C = M \log(P) + o(\log(P)), \quad (\text{A.7})$$

i.e., we have  $M$  DoF. Once again the DoF measure only the number of channels and not the channel strength or the noise power.

It is convenient to think of DoF as the number of signaling dimensions, where 1 signal dimension corresponds to one interference-free AWGN channel with SNR increasing proportionately with  $P$  as  $P$  approaches infinity. DoF are also known as the multiplexing gain since they measure the number of signals multiplexed over the air. Further, any signal carried by a wireless spectrum of two-sided baseband bandwidth  $B$  (one-sided bandwidth =  $B/2$ ) can be expressed by  $B$  freely chosen samples per sec according to the Nyquist–Shannon sampling theorem, and subject to the power constraints and the noise floor, each sample value can be seen as a signal dimension carrying 1 DoF. Thus, DoF can be equivalently understood as the bandwidth, multiplexing gain, number of signaling dimensions, or the capacity pre-log factor. As such, the fundamental significance of the DoF metric is evident.

# B

---

## Generalized Degrees of Freedom (GDoF)

---

An important limitation of the DoF metric is that it essentially forces all channels to be equally strong, i.e., the ratio of the signal powers (in dB scale) arriving on any two links approaches unity in the high SNR limit. It therefore does not provide much insight into the optimal ways to manage interference when some signals are significantly stronger or weaker than others. The *generalized* DoF metric serves precisely this purpose. The GDoF metric retains the analytical tractability of DoF as it explores the high SNR regime, but also fundamentally captures the diversity of signal strengths by fixing the ratios of different signal strengths (in dB scale) as all SNRs approach infinity. A potential difficulty with the GDoF metric is the explosion of parameters if all channels are assigned different relative strengths. Symmetric settings that are easier to visualize are therefore of particular interest, in spite of the obvious limitations imposed by the assumption of symmetry.

For example, consider the setting of Figure B.1 where the symmetric two user interference channel is illustrated. The interference channel has two independent messages  $W_1, W_2$  with message  $W_i$  originating at Transmitter  $i$  and intended for Receiver  $i$ ,  $i = 1, 2$ , respectively.

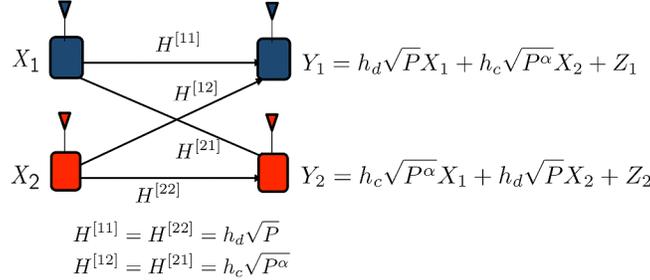


Fig. B.1 Two-user symmetric interference channel — GDoF setting.

Explicitly, the channel of Figure B.1 is defined by the input–output equations:

$$Y_1 = h_d\sqrt{P}X_1 + h_c\sqrt{P^\alpha}X_2 + Z_1 \quad (\text{B.1})$$

$$Y_2 = h_c\sqrt{P^\alpha}X_1 + h_d\sqrt{P}X_2 + Z_2 \quad (\text{B.2})$$

with the powers of the input symbols  $X_1, X_2$  and AWGN terms  $Z_1, Z_2$  normalized to unity. Similar to the DoF definition, the total GDoF metric,  $d(\alpha)$ , is defined as

$$d(\alpha) = \lim_{P \rightarrow \infty} \frac{C(P, \alpha)}{\log(P)}, \quad (\text{B.3})$$

where  $C(P, \alpha)$  is the sum-capacity parameterized by  $P, \alpha$ . Here  $\alpha$  is the ratio (in dB scale) of cross channel strength relative to direct channel strength. Note that the DoF metric corresponds to the point  $\alpha = 1$  within the GDoF characterization. Thus, the GDoF picture greatly broadens our understanding of optimal interference management techniques. It has been used most successfully to approximate the capacity of the two user interference channel to within a constant gap in [19, 40].

The GDoF curve of the symmetric two user interference channel of Figure B.1 is the highly recognizable “W” curve shown in Figure B.2. It greatly improves our understanding of the interference channel by identifying various regimes — corresponding to very weak, weak, moderately weak, strong and very strong interference — each of which has its own distinct character. Etkin et al. use the GDoF characterization as an insightful intermediate step leading to a capacity approximation of the two user interference channel that is accurate within one bit for all SNR

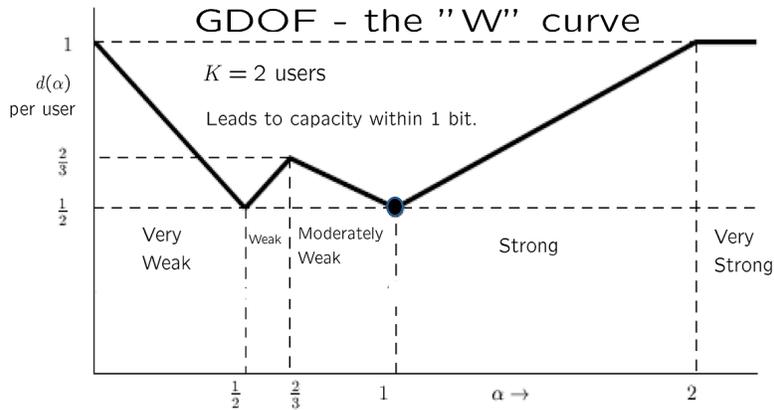


Fig. B.2 GDoF-per-user of the two-user symmetric interference channel — the “W” curve of Etkin et al. [40].

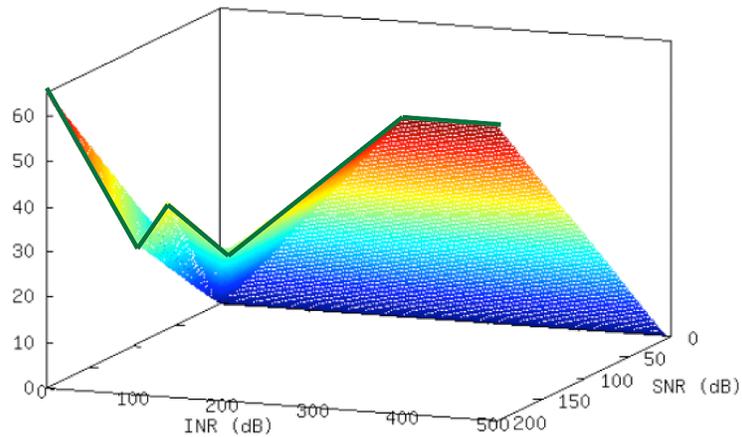


Fig. B.3 The manifestation of “W” curve. The plotted surface is a sum-rate capacity approximation (divided by the number of users = 2) that is accurate within a 1 bit gap.

values and all channel realizations. The appearance of the “W” curve in this 1 bit-gap approximation is illustrated in Figure B.3 where an achievable sum-rate (divided by the number of users) is plotted on the vertical axis as a function of signal to noise ratio,  $SNR = |h_d|^2 P$  and interference-to-noise ratio,  $INR = |h_c|^2 P^\alpha$ , both measured in dB scale. The surface representing sum-rate per user becomes an outer bound on achievable rates if elevated vertically by 1 bit, thus providing a capacity

approximation that is accurate within 1 bit. Note that no interference alignment is needed for this one-bit-gap capacity approximation.

Extensions of the two user interference channel GDoF to multiple antenna settings are explored in [75, 115], and the resulting GDoF curve for the 2 user MIMO interference channel where each transmitter is equipped with  $M$  antennas and each receiver is equipped with  $N$  antennas, is illustrated in Figure B.4. Constant gap capacity approximations are known in all cases of two user MIMO interference channel [74, 143]. Interestingly, the magnitude of the approximation gap scales in the number of antennas.

As another illustration of the GDoF metric, the natural generalization of the two user symmetric interference channel, is the  $(1 \times N)^{N+1}$  SIMO interference channel with  $N + 1$  users, where each transmitter has one antenna and each receiver has  $N$  antennas, all direct channels have strength  $P$  and all cross channels have strength  $P^\alpha$ . The GDoF of this setting are found by Gou and Jafar in [53] and are illustrated

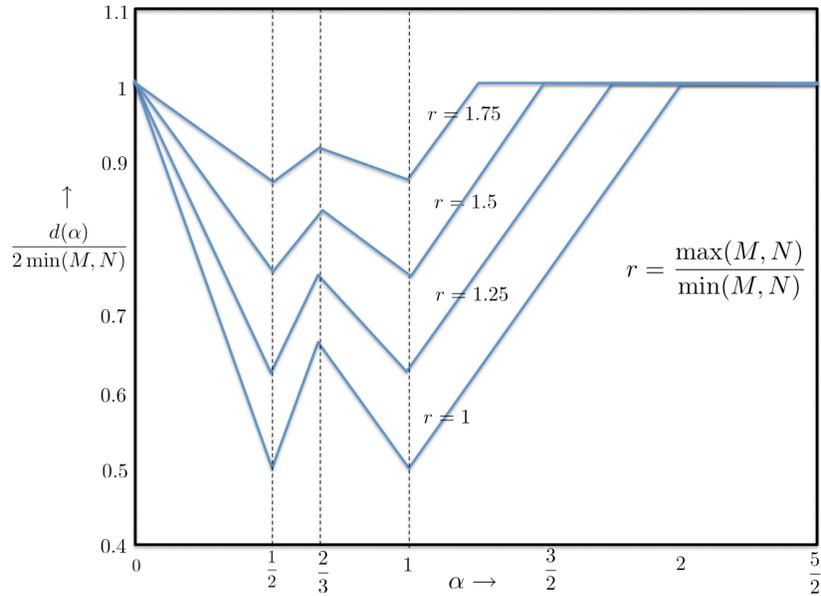


Fig. B.4 GDoF curves for the two-user MIMO interference channel with  $M$  antennas at each transmitter and  $N$  antennas at each receiver [75, 115].

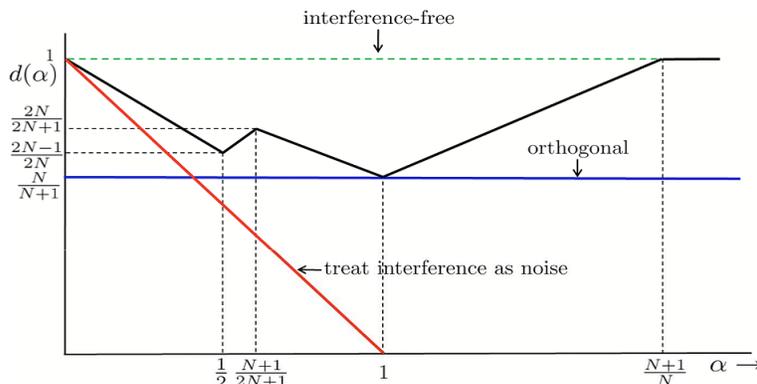


Fig. B.5 The GDoF-per-user curve for the SIMO  $(1 \times N)^{N+1}$  interference channel [53].

in Figure B.5. Clearly, setting  $N = 1$  in this curve produces the “W” curve of Figure B.2.

Just like the DoF characterizations, it is not only the GDoF value but also the new insights enabled by GDoF characterizations that make them particularly interesting. As an example, let us list some of the interesting observations from comparing the GDoF curve of the two user interference channel in Figure B.2 and the GDoF curve of its generalization, the  $(1 \times N)^{N+1}$  SIMO interference channel. Similar to the two user interference channel, the  $(1 \times N)^{N+1}$  SIMO interference channel does not require interference alignment to achieve the GDoF curve for  $N > 1$ . However, contrary to the two user interference channel where treating interference as noise (i.e., sending only private messages) is GDoF-optimal in the very weak interference regime, the  $(1 \times N)^{N+1}$  SIMO interference channel does not possess a regime where treating interference as noise is GDoF-optimal. The sub-optimality of treating interference as noise is indicated in Figure B.5 by the strict dominance of the GDoF curve over the line representing the achievable GDoF with interference treated as noise. Also interesting is the observation that the layered deterministic model of Avestimehr et al. [3], which is the key to the GDoF characterization of the two user interference channel found by Etkin et al. in [40], does not extend sufficiently to the  $(1 \times N)^{N+1}$  SIMO interference channel setting. Gou and Jafar instead rely on an extension of an earlier deterministic interference channel

capacity result by El Gamal and Costa in [42] for their GDoF result. It is interesting to note that the  $(1 \times N)^{N+1}$  SIMO interference channel ends up being the natural setting to exploit an extension of the two user deterministic interference channel setting of El Gamal and Costa, mainly because the channel does not allow interference alignment in the GDoF sense. Intuitively, this is because in the SIMO interference channel, after decoding its intended message and subtracting the decoded codeword from the received signal, any receiver has  $N$  receive antennas which provide the receiver  $N$  linearly independent equations in the  $N$  interfering symbols. Thus, the interference dimensions remain resolvable, i.e., cannot be aligned. Note that this is also, intuitively, the characteristic of the deterministic interference channel setting studied by El Gamal and Costa in [42], i.e., the interfering signal should be resolvable from the received signal at any receiver, given the desired signal. The  $(1 \times N)^{N+1}$  SIMO interference channel setting is special in this regard. On the one hand, if we decrease the number of users, say  $(1 \times N)^N$ , the GDoF characterization will become trivial, because simple zero forcing will allow each user to access 1 DoF, i.e., its maximum interference-free value, regardless of the strength of interference. This is because each receiver will have enough antennas to remove all interference. On the other hand if we increase the number of users, say  $(1 \times N)^{N+2}$ , then interference alignment will play a central role in the GDoF characterization, thereby fundamentally changing the character of the problem. Finally, while the GDoF characterization for the two user interference channel directly leads to a 1 bit gap capacity approximation regardless of channel realizations and SNR, INR values, the GDoF characterization of the  $(1 \times N)^{N+1}$  SIMO interference channel only produces a bounded gap capacity characterization, specifically an  $O(1)$  approximation, where the gap does not depend on SNR and INR values (for almost all channel realizations) but the magnitude of the gap depends on the channel realizations in general. This is easy to see,

because the  $(1 \times N)^{N+1}$  SIMO interference channel, for some channel realizations, devolves into the SISO interference channel with more than two users for which interference alignment is very much a factor in determining the capacity pre-log and no bounded gap capacity characterizations are known.

## References

---

- [1] M. Aldridge, O. Johnson, and R. Piechocki, "Asymptotic sum-capacity of random Gaussian interference networks using interference alignment," in *Proceedings of IEEE International Symposium on Information Theory (ISIT), 2010*, pp. 410–414, IEEE, 2010.
- [2] V. Annapureddy, A. El Gamal, and V. Veeravalli, "Degrees of freedom of the K-user interference channel with transmitter cooperation," in *IEEE International Symposium on Information Theory Proceedings (ISIT) 2010*, pp. 385–389, IEEE, 2010.
- [3] A. S. Avestimehr, S. Diggavi, and D. Tse, "Wireless Network Information Flow: A Deterministic Approach," *IEEE Transactions on Information Theory*, vol. 57, pp. 1872–1905, 2011.
- [4] S. Avestimehr, A. Sezgin, and D. Tse, "Capacity of the Two Way Relay Channel within a Constant Gap," *European Transactions on Telecommunications*, vol. 21, pp. 363–374, April 2010.
- [5] O. E. Ayach and R. W. Heath, Jr., "Interference alignment with analog channel state feedback," *CoRR*, vol. abs/1010.2787, [Online]. Available: <http://arxiv.org/abs/1010.2787>, 2010.
- [6] G. Bagherikaram, A. S. Motahari, and A. K. Khandani, "On the secure degrees-of-freedom of the multiple-access-channel," *CoRR*, vol. abs/1003.0729, [Online]. Available: <http://arxiv.org/abs/1003.0729>, 2010.
- [7] B. Bandemer, A. E. Gamal, and G. Vazquez-Vilar, "On the sum capacity of a class of cyclically symmetric deterministic interference channels," in *IEEE International Symposium on Information Theory, ISIT*, pp. 2622–2626, 2009.
- [8] R. Bassily and S. Ulukus, "Ergodic secret alignment," *CoRR*, vol. abs/1010.6057, [Online]. Available: <http://arxiv.org/abs/1010.6057>, 2010.

- [9] T. D. Bavarisetti, G. Abhinav, K. Prasad, and B. S. Rajan, "A transform approach to linear network coding for acyclic networks with delay," *CoRR*, vol. abs/1103.3882, [Online]. Available: <http://arxiv.org/abs/1103.3882>, 2011.
- [10] V. Beresnevich, "A Groshev type theorem for convergence on manifolds," *Acta Mathematica Hungarica* 94, vol. 1–2, pp. 99–130, 2002.
- [11] S. Berger, T. Unger, M. Kuhn, A. Klein, and A. Wittneben, "Recent advances in amplify-and-forward two-hop relaying," *Communications Magazine*, [Online]. Available: <http://www.nari.ee.ethz.ch/wireless/pubs/p/CommMag08>, July 2009.
- [12] I. Bernik, D. Kleinbock, and G. Margulis, "Khintchine-type theorems on manifolds: The convergence case for standard and multiplicative versions," *International Mathematics Research Notices*, vol. 9, pp. 453–386, 2001.
- [13] Y. Birk and T. Kol, "Informed-source coding-on-demand (ISCOD) over broadcast channels," in *Proceedings of the Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies, IEEE INFOCOM'98*, pp. 1257–1264, 1998.
- [14] Y. Birk and T. Kol, "Coding on demand by an informed source (ISCOD) for efficient broadcast of different supplemental data to caching clients," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2825–2830, June 2006.
- [15] H. Boelcskei, R. Nabar, O. Oyman, and A. Paulraj, "Capacity scaling laws in MIMO relay networks," *Transactions on Wireless Communications*, vol. 5, no. 6, pp. 1433–1444, June 2006.
- [16] S. Borade, L. Zheng, and R. Gallager, "Maximizing degrees of freedom in wireless networks," in *Proceedings of 40th Annual Allerton Conference on Communication, Control and Computing*, pp. 561–570, October 2003.
- [17] G. Bresler, D. Cartwright, and D. Tse, "Settling the feasibility of interference alignment for the MIMO interference channel: The symmetric square case," *CoRR*, vol. abs/1104.0888, [Online]. Available: <http://arxiv.org/abs/1104.0888>, 2011.
- [18] G. Bresler, A. Parekh, and D. Tse, "The approximate capacity of the many-to-one and one-to-many Gaussian interference channels," *IEEE Transactions on Information Theory*, vol. 56, pp. 4566–4592, September 2010.
- [19] G. Bresler and D. Tse, "The two-user Gaussian interference channel: A deterministic view," *European Transactions in Telecommunications*, vol. 19, no. 4, pp. 333–354, June 2008.
- [20] G. Bresler and D. Tse, "3 user interference channel: Degrees of freedom as a function of channel diversity," in *47th Annual Allerton Conference on Communication, Control, and Computing*, pp. 265–271, 2009.
- [21] V. Cadambe, C. Huang, S. A. Jafar, and J. Li, "Optimal repair of MDS codes in distributed storage via subspace interference alignment," *Arxiv preprint arXiv:1106.1250*, [Online]. Available: <http://arxiv.org/pdf/1106.1250v1>, 2011.
- [22] V. Cadambe and S. Jafar, "Interference alignment and the degrees of freedom of the K user interference channel," *IEEE Transactions on Information Theory*, vol. 54, no. 8, pp. 3425–3441, August 2008.

- [23] V. Cadambe and S. Jafar, "Interference alignment and the degrees of freedom of wireless X networks," *IEEE Transactions on Information Theory*, no. 9, pp. 3893–3908, September 2009.
- [24] V. Cadambe and S. Jafar, "Parallel Gaussian interference channels are not always separable," *IEEE Transactions on Information Theory*, vol. 55, no. 9, pp. 3983–3990, 2009.
- [25] V. Cadambe and S. Jafar, "Sum-capacity and the unique separability of the parallel Gaussian MAC-Z-BC network," in *Proceedings of IEEE International Symposium on Information Theory (ISIT), 2010*, pp. 2318–2322, 2010.
- [26] V. Cadambe, S. Jafar, and S. Shamai, "Interference alignment on the deterministic channel and application to Gaussian networks," *IEEE Transactions on Information Theory*, vol. 55, no. 1, pp. 269–274, January 2009.
- [27] V. Cadambe, S. Jafar, and C. Wang, "Interference alignment with asymmetric complex signaling settling the Høst-Madsen–Nosratinia conjecture," *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4552–4565, 2010.
- [28] V. R. Cadambe and S. A. Jafar, "Degrees of freedom of wireless networks with relays, feedback, cooperation and full duplex operation," *IEEE Transactions on Information Theory*, vol. 55, pp. 2334–2344, May 2009.
- [29] V. R. Cadambe, S. A. Jafar, and H. Maleki, "Distributed data storage with minimum storage regenerating codes — exact and functional repair are asymptotically equally efficient," *CoRR*, vol. abs/1004.4299, [Online]. Available: <http://arxiv.org/abs/1004.4299>, 2010.
- [30] K. Cai, K. Letaief, P. Fan, and R. Feng, "On the solvability of 2-pair unicast networks—A cut-based characterization," *Arxiv preprint arXiv:1007.0465*, 2010.
- [31] G. Caire and S. Shamai, "On the achievable throughput of a multiantenna Gaussian broadcast channel," *IEEE Transactions on Information Theory*, vol. 49, no. 7, pp. 1691–1706, July 2003.
- [32] Y. Cao and B. Chen, "Capacity bounds for two-hop interference networks," *47th Annual Allerton Conference on Communication, Control, and Computing*, vol. abs/0910.1532, [Online]. Available: <http://arxiv.org/abs/0910.1532>, 2009.
- [33] S. Chen and R. Cheng, "Achieve the degrees of freedom of K-User MIMO interference channel with a MIMO relay," in *IEEE GLOBECOM*, December 2010.
- [34] S. W. Choi and S.-Y. Chung, "On the separability of parallel Gaussian interference channels," *Proceedings of IEEE International Symposium on Information Theory (ISIT 2009)*, pp. 2592–2596, [Online]. Available: <http://arxiv.org/abs/0905.1537>, 2009.
- [35] B. Da and R. Zhang, "Exploiting interference alignment in multi-cell cooperative OFDMA resource allocation," *CoRR*, vol. abs/1103.3093, [Online]. Available: <http://arxiv.org/abs/1103.3093>, 2011.
- [36] A. K. Das, S. Vishwanath, S. A. Jafar, and A. Markopoulou, "Network coding for multiple unicasts: An interference alignment approach," *CoRR*, vol. abs/1008.0235, [Online]. Available: <http://arxiv.org/abs/1008.0235>, 2010.

- [37] A. G. Dimakis, K. Ramchandran, Y. Wu, and C. Suh, "A survey on network codes for distributed storage," *Proceedings of the IEEE*, vol. 99, pp. 476–489, [Online]. Available: <http://arxiv.org/abs/1004.4438>, 2011.
- [38] C. Esli and A. Wittneben, "A hierarchical AF protocol for distributed orthogonalization in multiuser relay networks," *IEEE Transactions on Vehicular Technology*, vol. 59, pp. 3902–3916, [Online]. Available: [http://www.nari.ee.ethz.ch/wireless/pubs/p/esli\\_tvt\\_2010](http://www.nari.ee.ethz.ch/wireless/pubs/p/esli_tvt_2010), August 2010.
- [39] R. Etkin and E. Ordentlich, "The degrees-of-freedom of the K-User Gaussian interference channel is discontinuous at rational channel coefficients," *IEEE Transactions on Information Theory*, vol. 55, pp. 4932–4946, November 2009.
- [40] R. Etkin, D. Tse, and H. Wang, "Gaussian interference channel capacity to within one bit," *IEEE Transactions on Information Theory*, vol. 54, no. 12, pp. 5534–5562, 2008.
- [41] G. J. Foschini and M. J. Gans, "On limits of wireless communications in a fading environment when using multiple antennas," *Wireless Personal Commun. : Kluwer Academic Press*, no. 6, pp. 311–335, 1998.
- [42] A. E. Gamal and M. Costa, "The capacity region of a class of deterministic interference channels," *IEEE Transactions on Information Theory*, vol. 2, pp. 343–346, March 1982.
- [43] D. Gesbert, S. Hanly, H. Huang, S. Shitz, O. Simeone, and W. Yu, "Multi-cell MIMO cooperative networks: A new look at interference," *IEEE Journal on Selected Areas in Communications*, vol. 28, no. 9, pp. 1380–1408, December 2010.
- [44] A. Ghasemi, A. Motahari, and A. Khandani, "Interference alignment for the K user MIMO interference channel," in *Proceedings of International Symposium on Information Theory (ISIT)*, pp. 360–364, IEEE, 2010.
- [45] A. Goldsmith, S. A. Jafar, N. Jindal, and S. Vishwanath, "Capacity limits of MIMO channels," *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 5, pp. 684–702, June 2003.
- [46] A. Goldsmith and P. Varaiya, "Capacity of fading channels with channel side information," *IEEE Transactions on Information Theory*, vol. 43, pp. 1986–1992, November 1997.
- [47] S. Gollakota, S. Perli, and D. Katabi, "Interference alignment and cancellation," in *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, pp. 159–170, ACM, 2009.
- [48] K. Gomadam, V. Cadambe, and S. Jafar, "A distributed numerical approach to interference alignment and applications to wireless interference networks," *IEEE Transactions on Information Theory*, pp. 3309–3322, June 2011.
- [49] T. Gou and S. Jafar, "On the secure degrees of freedom of wireless X networks," in *Proceedings of 46th Annual Allerton Conference on Communication, Control and Computing*, September 2008.
- [50] T. Gou and S. Jafar, "Degrees of freedom of the K user  $M \times N$  MIMO interference channel," *IEEE Transactions on Information Theory*, vol. 56, no. 12, pp. 6040–6057, December 2010.
- [51] T. Gou, S. Jafar, S. Jeon, and S. Chung, "Aligned interference neutralization and the degrees of freedom of the  $2 \times 2 \times 2$  interference channel," *Proceedings*

- of *IEEE International Symposium on Information Theory (ISIT)*, [Online]. Available: <http://arxiv.org/abs/1012.2350>, 2011.
- [52] T. Gou, S. Jafar, and C. Wang, “Degrees of freedom of finite state compound wireless networks,” *To Appear in the IEEE Transactions on Information Theory*. Full paper available at *arXiv:0909.4203*, [Online]. Available: <http://arxiv.org/abs/0909.5424>, 2009.
- [53] T. Gou and S. A. Jafar, “Capacity of a class of symmetric SIMO Gaussian interference channels within  $O(1)$ ,” *IEEE Transactions on Information Theory*, vol. 57, no. 4, pp. 1932–1958, April 2011.
- [54] T. Gou, C. Wang, and S. A. Jafar, “Aiming perfectly in the dark — Blind interference alignment through staggered antenna switching,” *IEEE Transactions on Signal Processing*, vol. 59, pp. 2734–2744, June 2011.
- [55] L. Gropop, D. Tse, and R. Yates, “Interference alignment for line-of-sight channels,” *CORR*, vol. 0809.3035, [online]. Available: <http://arXiv:cs.IT/0809.3035>, 2008.
- [56] X. He and A. Yener, “K-user interference channels: Achievable secrecy rate and degrees of freedom,” *CoRR*, vol. abs/0905.2643, [Online]. Available: <http://arxiv.org/abs/0905.2643>, 2009.
- [57] A. Host-Madsen and A. Nosratinia, “The multiplexing gain of wireless networks,” in *Proceedings of ISIT*, 2005.
- [58] C. Huang, V. Cadambe, and S. Jafar, “On the capacity and generalized degrees of freedom of the X channel,” *arxiv:0810.4741*, October 2008.
- [59] C. Huang, S. A. Jafar, S. Shamai, and S. Vishwanath, “On degrees of freedom region of MIMO networks without CSIT,” *CoRR*, vol. abs/0909.4017, [Online]. Available: <http://arxiv.org/abs/0909.4017>, 2009.
- [60] H. Huang, V. Lau, Y. Du, and S. Liu, “Robust lattice alignment for K-user MIMO interference channels with imperfect channel knowledge,” *IEEE Transactions on Signal Processing*, vol. 59, pp. 3315–3325, 2011.
- [61] S. Jafar, “Too much mobility limits the capacity of wireless ad-hoc networks,” *IEEE Transactions on Information Theory*, vol. 52, no. 7, pp. 3954–3965, November 2005.
- [62] S. Jafar, “Degrees of Freedom on the MIMO X channel- optimality of the MMK scheme,” *CoRR*, vol. abs/cs/0607099v2, [Online]. Available: <http://arxiv.org/abs/cs/0607099v2>, September 2006.
- [63] S. Jafar, “Exploiting channel correlations — Simple interference alignment schemes with no CSIT,” *Preprint available on Arxiv arXiv:0910.0555*, October 2009.
- [64] S. Jafar and V. Cadambe, “Degrees of freedom of wireless networks — what a difference delay makes,” in *Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA*, November 2007.
- [65] S. Jafar and M. Fakhreddin, “Degrees of freedom for the MIMO interference channel,” *IEEE Transactions on Information Theory*, vol. 53, no. 7, pp. 2637–2642, July 2007.
- [66] S. Jafar and A. Goldsmith, “Isotropic fading vector broadcast channels: the scalar upperbound and loss in degrees of freedom,” *IEEE Transactions on Information Theory*, vol. 51, no. 3, pp. 848–857, March 2005.

- [67] S. Jafar and S. Shamai, “Degrees of freedom region for the MIMO X channel,” *IEEE Transactions on Information Theory*, vol. 54, no. 1, pp. 151–170, January 2008.
- [68] S. Jafar and S. Vishwanath, “Generalized degrees of freedom of the symmetric gaussian K user interference channel,” *IEEE Transactions on Information Theory*, vol. 56, no. 7, pp. 3297–3303, July 2010.
- [69] S. A. Jafar, “The ergodic capacity of interference networks,” *CoRR*, vol. abs/0902.0838, [Online]. Available: <http://arxiv.org/abs/0902.0838>, 2009.
- [70] S.-W. Jeon and S.-Y. Chung, “Capacity of a class of multi-source relay networks,” *CoRR*, vol. abs/0907.2510, [Online]. Available: <http://arxiv.org/abs/0907.2510>, 2009.
- [71] S.-W. Jeon, S.-Y. Chung, and S. Jafar, “Degrees of freedom region of a class of multisource Gaussian relay networks,” *IEEE Transactions on Information Theory*, vol. 57, pp. 3032–3044, May 2011.
- [72] O. Johnson, M. Aldridge, and R. Piechocki, “Interference alignment-based sum capacity bounds for random dense Gaussian interference networks,” *IEEE Transactions on Information Theory*, vol. 57, pp. 282–290, [Online]. Available: <http://arxiv.org/abs/1004.0208>, January 2011.
- [73] O. Johnson, M. Aldridge, and R. J. Piechocki, “Delay-rate tradeoff in ergodic interference alignment,” *CoRR*, vol. abs/1004.0208, [Online]. Available: <http://arxiv.org/abs/1004.0208>, 2010.
- [74] S. Karmakar and M. K. Varanasi, “Capacity of the MIMO interference channel to within a constant gap,” *CoRR*, vol. abs/1102.0267, [Online]. Available: <http://arxiv.org/abs/1102.0267>, 2011.
- [75] S. Karmakar and M. K. Varanasi, “The Generalized Degrees of Freedom Region of the MIMO Interference Channel,” *CoRR*, vol. abs/1103.2560, [Online]. Available: <http://arxiv.org/abs/1103.2560>, 2011.
- [76] S. Katti, S. Gollakota, and D. Katabi, “Embracing wireless interference: Analog network coding,” *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4, pp. 397–408, 2007.
- [77] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft, “XORs in the air: practical wireless network coding,” *IEEE/ACM Transactions on Networking (TON)*, vol. 16, no. 3, pp. 497–510, 2008.
- [78] L. Ke, A. Ramamoorthy, Z. Wang, and H. Yin, “Degrees of Freedom Region for an Interference Network with General Message Demands,” *CoRR*, vol. abs/1101.3068, [Online]. Available: <http://arxiv.org/abs/1101.3068>, 2011.
- [79] L. Ke and Z. Wang, “Degrees of Freedom Regions of Two-User MIMO Z and Full Interference Channels: The Benefit of Reconfigurable Antennas,” *ArXiv e-prints*, November 2010.
- [80] A. Khisti, “Interference Alignment for the Multi-Antenna Compound Wiretap Channel,” *IEEE Transactions on Information Theory*, vol. 57, pp. 2976–2993, May 2011.
- [81] J. C. Koo, W. Wu, and J. T. G. III, “Delay-rate tradeoff for ergodic interference alignment in the Gaussian case,” *48th Annual Allerton Conference on Communication, Control, and Computing*, vol. abs/1001.2582, pp. 1069–1075, [Online]. Available: <http://arxiv.org/abs/1001.2582>, 2010.

- [82] J. Korner and K. Marton, "How to encode the modulo-two sum of binary sources," *IEEE Transactions on Information Theory*, vol. 25, pp. 219–221, March 1979.
- [83] O. O. Koyluoglu, H. E. Gamal, L. Lai, and H. V. Poor, "Interference Alignment for Secrecy," *IEEE Transactions on Information Theory*, vol. 57, pp. 3323–3332, June 2011.
- [84] G. Kramer, I. Marić, and R. Yates, "Cooperative communications," *Foundations and Trends® in Networking*, vol. 1, no. 3, pp. 271–425, 2006.
- [85] R. T. Krishnamachari and M. K. Varanasi, "Interference Alignment Under Limited Feedback for MIMO Interference Channels," *CoRR*, vol. abs/0911.5509, [Online]. Available: <http://arxiv.org/abs/0911.5509>, 2009.
- [86] A. Lapidoth, S. Shamai, and M. Wigger, "On the capacity of Fading MIMO broadcast channels with imperfect transmitter side-information," in *Proceedings of 43rd Annual Allerton Conference on Communications, Control and Computing*, pp. 28–30, September 2005.
- [87] N. Lee, J. Lim, and J. Chun, "Degrees of freedom of the MIMO Y channel: signal space alignment for network coding," *IEEE Transactions on Information Theory*, vol. 56, no. 7, pp. 3332–3342, 2010.
- [88] N. Lee, W. Shin, and B. Clerckx, "Interference Alignment with Limited Feedback on Two-cell Interfering Two-User MIMO-MAC," *CoRR*, vol. abs/1010.0933, [Online]. Available: <http://arxiv.org/abs/1010.0933>, 2010.
- [89] M. Maddah-Ali, A. Motahari, and A. Khandani, "Communication over X channel: Signaling and Performance Analysis," in *Technical Report. UW-ECE-2006-27, University of Waterloo*, December 2006.
- [90] M. Maddah-Ali, A. Motahari, and A. Khandani, "Communication over X channel: Signalling and multiplexing gain," in *Technical Report. UW-ECE-2006-12, University of Waterloo*, July 2006.
- [91] M. Maddah-Ali, A. Motahari, and A. Khandani, "Signaling over MIMO multi-base systems - combination of multi-access and broadcast schemes," in *Proceedings of ISIT*, pp. 2104–2108, July 2006.
- [92] M. Maddah-Ali, A. Motahari, and A. Khandani, "Communication over MIMO X Channels: Interference Alignment, Decomposition, and Performance Analysis," in *IEEE Transactions on Information Theory*, pp. 3457–3470, August 2008.
- [93] M. A. Maddah-Ali, "The Degrees of Freedom of the Compound MIMO Broadcast Channels with Finite States," *CoRR*, vol. abs/0909.5006, [Online]. Available: <http://arxiv.org/abs/0909.5006>, 2009.
- [94] M. A. Maddah-Ali and D. Tse, "On the Degrees of Freedom of MISO Broadcast Channels with Delayed Feedback," Technical Report UCB/EECS-2010-122, EECS Department, University of California, Berkeley, [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-122.html>, September 2010.
- [95] H. Maleki, S. A. Jafar, and S. Shamai, "Retrospective Interference Alignment," *CoRR*, vol. abs/1009.3593, [Online]. Available: <http://arxiv.org/abs/1009.3593>, 2010.

- [96] R. Mathar and M. Zivkovic, "How to position  $n$  transmitter-receiver pairs in  $n-1$  dimensions such that each can use half of the channel with zero interference from the others," in *Global Telecommunications Conference, 2009. GLOBECOM 2009. IEEE*, pp. 1–4, IEEE, 2010.
- [97] S. Mohajer, S. Diggavi, C. Fragouli, and D. Tse, "Approximate Capacity of a Class of Gaussian Interference-Relay Networks," *IEEE Transactions on Information Theory*, vol. 57, pp. 2837–2864, May 2011.
- [98] V. Morgenshtern and H. Bölcskei, "Crystallization in large wireless networks," *IEEE Transactions on Information Theory*, vol. 53, no. 10, pp. 3319–3349, [Online]. Available: <http://www.nari.ee.ethz.ch/comwth/pubs/p/transit06>, October 2007.
- [99] A. Motahari, S. Gharan, M. Maddah-Ali, and A. Khandani, "Real Interference Alignment: Exploiting the Potential of Single Antenna Systems," *CoRR*, vol. abs/0908.2282, [Online]. Available: <http://arxiv.org/abs/0908.2282>, 2009.
- [100] B. Nazer and M. Gastpar, "Compute-and-Forward: Harnessing Interference through Structured Codes," *CoRR*, vol. abs/0908.2119, [Online]. Available: <http://arxiv.org/abs/0908.2119>, 2009.
- [101] B. Nazer, M. Gastpar, S. Jafar, and S. Vishwanath, "Ergodic Interference Alignment," in *ISIT*, 2009.
- [102] B. Nazer, M. Gastpar, S. Jafar, and S. Vishwanath, "Interference alignment at finite SNR: General message sets," in *47th Annual Allerton Conference on Communication, Control, and Computing*, pp. 843–848, IEEE, 2009.
- [103] B. Nazer, A. Sanderovich, M. Gastpar, and S. Shamai, "Structured Superposition for Backhaul Constrained Cellular Uplink," in *IEEE ISIT 2009*, June–July 2009.
- [104] F. Negro, S. Shenoy, D. Slock, and I. Ghauri, "Interference alignment limits for  $K$ -user frequency-flat MIMO interference channels," in *Proceedings European Signal Proc. Conf. (Eusipco)*, 2009.
- [105] U. Niesen, "Interference alignment in dense wireless networks," *IEEE Transactions on Information Theory*, vol. 57, pp. 2889–2901, [Online]. Available: <http://arxiv.org/abs/0912.0868>, May 2011.
- [106] U. Niesen and P. Whiting, "The Degrees of Freedom of Compute-and-Forward," *CoRR*, vol. abs/1101.2182, [Online]. Available: <http://arxiv.org/abs/1101.2182>, 2011.
- [107] H. Ning, C. Ling, and K. Leung, "Relay-aided interference alignment: Feasibility conditions and algorithm," in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, pp. 390–394, IEEE, 2010.
- [108] B. Niu and A. Haimovich, "Interference subspace tracking for network interference alignment in cellular systems," in *Global Telecommunications Conference, 2009. GLOBECOM 2009. IEEE*, pp. 1–5, IEEE, 2010.
- [109] B. Nosrat-Makouei, J. G. Andrews, and R. W. H. Jr., "MIMO Interference Alignment Over Correlated Channels with Imperfect CSI," *CoRR*, vol. abs/1010.2741, [Online]. Available: <http://arxiv.org/abs/1010.2741>, 2010.
- [110] B. Nourani, S. Motahari, and A. Khandani, "Relay-aided interference alignment for the quasi-static X channel," in *IEEE International Symposium on Information Theory, ISIT*, pp. 1764–1768, IEEE, 2009.

- [111] B. Nourani, S. Motahari, and A. Khandani, "Relay-aided Interference Alignment for the quasi-static interference channel," in *IEEE International Symposium on Information Theory Proceedings (ISIT)*, pp. 405–409, IEEE, 2010.
- [112] D. Papailiopoulos and A. Dimakis, "Interference alignment as a rank constrained rank minimization," *IEEE GLOBECOM*, December 2010.
- [113] D. Papailiopoulos and A. Dimakis, "Distributed Storage Codes through Hadamard Designs," *Arxiv preprint arXiv:1106.1652*, [Online]. Available: <http://arxiv.org/pdf/1106.1652>, 2011.
- [114] J. Park, Y. Sung, and V. Poor, "On Beamformer Design for Multiuser MIMO Channels," *CoRR*, vol. abs/1011.6121, [Online]. Available: <http://arxiv.org/abs/1011.6121>, November 2010.
- [115] P. Parker, D. Bliss, and V. Tarokh, "On the Degrees-of-Freedom of the MIMO Interference Channel," in *42nd Annual Conference on Information Sciences and Systems (CISS)*, March 2008.
- [116] S. Peters and R. Heath, "Interference alignment via alternating minimization," in *IEEE International Conference on Acoustics, Speech and Signal Processing, 2009. ICASSP 2009*, pp. 2445–2448, IEEE, 2009.
- [117] S. W. Peters and R. W. Heath, Jr., "User partitioning for less overhead in MIMO interference channels," *CoRR*, vol. abs/1007.0512, [Online]. Available: <http://arxiv.org/abs/1007.0512>, 2010.
- [118] S. W. Peters and R. W. Heath Jr., "Cooperative Algorithms for MIMO Interference Channels," *IEEE Transactions on Vehicular Technology*, vol. 60, pp. 206–218, January 2011.
- [119] T. Philosof and R. Zamir, "On the loss of single-letter characterization: the dirty multiple access channel," *IEEE Transactions on Information Theory*, vol. 55, no. 6, pp. 2442–2454, 2009.
- [120] A. Ramakrishnan, A. Das, H. Maleki, A. Markopoulou, S. Jafar, and S. Vishwanath, "Network coding for three unicast sessions: Interference alignment approaches," *Allerton Conference on Communications, Control and Computing*, October 2010.
- [121] B. Rankov and A. Wittneben, "Achievable rate regions for the two-way relay channel," in *Proceedings IEEE International Symposium on Information Theory (ISIT)*, July 2006. [Online]. Available: <http://www.nari.ee.ethz.ch/wireless/pubs/p/isit2006>.
- [122] B. Rankov and A. Wittneben, "Spectral efficient protocols for half-duplex fading relay channels," *IEEE Journal on Selected Areas in Communications*, [Online]. Available: <http://www.nari.ee.ethz.ch/wireless/pubs/p/jsac2006>, February 2007.
- [123] M. Razaviyayn, G. Lyubeznik, and Z. Luo, "On the degrees of freedom achievable through interference alignment in a MIMO interference channel," *CoRR*, vol. abs/1104.0992, [Online]. Available: <http://arxiv.org/abs/1104.0992>, 2011.
- [124] M. Razaviyayn, M. Sanjabi, and Z.-Q. Luo, "Linear Transceiver Design for Interference Alignment: Complexity and Computation," *2010 IEEE Eleventh International Workshop on Signal Processing Advances in*

- Wireless Communications (SPAWC)*, vol. abs/1009.3481, [Online]. Available: <http://arxiv.org/abs/1009.3481>, 2010.
- [125] S. Saha and R. Berry, "On the Sum Capacity of a Class of 3 User Deterministic Interference Channels," in *Forty-Eighth Annual Allerton Conference on Communication, Control and Computing*, September 2010.
- [126] L. Sankar, X. Shang, E. Erkip, and H. Poor, "Ergodic Fading Interference Channels: Sum-Capacity and Separability," *IEEE Transactions on Information Theory*, vol. 57, pp. 2605–2626, May 2011.
- [127] D. Schmidt, C. Shi, R. Berry, M. Honig, and W. Utschick, "Minimum mean squared error interference alignment," in *Conference Record of the Forty-Third Asilomar Conference on Signals, Systems and Computers, 2009*, pp. 1106–1110, IEEE, 2010.
- [128] D. Schmidt, W. Utschick, and M. Honig, "Large System Performance of Interference Alignment in Single-Beam MIMO Networks," in *IEEE Global Telecommunications Conference (GLOBECOM) 2010*, December 2010.
- [129] A. Sezgin, A. S. Avestimehr, M. A. Khajehnejad, and B. Hassibi, "Divide-and-conquer: Approaching the capacity of the two-pair bidirectional Gaussian relay network," *CoRR*, vol. abs/1001.4271, [Online]. Available: <http://arxiv.org/abs/1001.4271>, 2010.
- [130] N. B. Shah, K. V. Rashmi, P. V. Kumar, and K. Ramchandran, "Explicit Codes Minimizing Repair Bandwidth for Distributed Storage," *CoRR*, vol. abs/0908.2984, [Online]. Available: <http://arxiv.org/abs/0908.2984>, 2009.
- [131] H. Shen, B. Li, M. Tao, and Y. Luo, "The new interference alignment scheme for the MIMO interference channel," in *Wireless Communications and Networking Conference (WCNC), 2010 IEEE*, pp. 1–6, IEEE, 2010.
- [132] S. Shenvi and B. Dey, "A simple necessary and sufficient condition for the double unicast problem," in *IEEE International Conference on Communications (ICC), 2010*, pp. 1–5, IEEE, 2010.
- [133] W. Shin, N. Lee, J. B. Lim, C. Shin, and K. Jang, "Interference alignment through user cooperation for two-cell MIMO interfering broadcast channels," *CoRR*, vol. abs/1011.3867, [Online]. Available: <http://arxiv.org/abs/1011.3867>, 2010.
- [134] I. Shomorony and S. Avestimehr, "Two-Unicast Wireless Networks: Characterizing the Sum Degrees of Freedom," *CoRR*, vol. abs/1102.2498, [Online]. Available: <http://arxiv.org/abs/1102.2498>, 2011.
- [135] O. Simeone, O. Somekh, Y. Bar-Ness, H. V. Poor, and S. Shamai, "Capacity of Linear Two-hop Mesh Networks with Rate Splitting, Decode-and-forward Relaying and Cooperation," *proceedings of the 45th Annual Allerton Conference on Communication, Control and Computing, Monticello, IL*, vol. abs/0710.2553, [Online]. Available: <http://arxiv.org/abs/0710.2553>, September 2007.
- [136] S. Sridharan, A. Jafarian, S. Vishwanath, and S. Jafar, "Capacity of Symmetric K-User Gaussian Very Strong Interference Channels," in *Proceedings of IEEE GLOBECOM*, December 2008.

- [137] C. Suh and K. Ramchandran, "Exact Regeneration Codes for Distributed Storage Repair Using Interference Alignment," *CoRR*, vol. abs/1001.0107, [Online]. Available: <http://arxiv.org/abs/1001.0107>, 2010.
- [138] C. Suh and K. Ramchandran, "On the Existence of Optimal Exact-Repair MDS Codes for Distributed Storage," *CoRR*, vol. abs/1004.4663, [Online]. Available: <http://arxiv.org/abs/1004.4663>, 2010.
- [139] C. Suh and D. Tse, "Interference Alignment for Cellular Networks," in *Proceedings of 40th Annual Allerton Conference on Communication, Control and Computing*, September 2008.
- [140] I. Tamo, Z. Wang, and J. Bruck, "MDS array codes with optimal rebuilding," in *IEEE ISIT 2011*, July–August 2011.
- [141] R. Tannious and A. Nosratinia, "The interference channel with MIMO relay: Degrees of freedom," in *IEEE International Symposium on Information Theory, (ISIT) 2008*, pp. 1908–1912, IEEE, 2008.
- [142] E. Telatar, "Capacity of Multi-antenna Gaussian Channels," *European Transactions on Telecomm. ETT*, vol. 10, no. 6, pp. 585–596, November 1999.
- [143] E. Telatar and D. Tse, "Bounds on the capacity region of a class of interference channels," in *Proceedings of IEEE International Symposium on Information Theory (ISIT)*, 2007.
- [144] J. Thukral and H. Boelcskei, "Interference alignment with limited feedback," in *Proceedings of the 2009 IEEE International Symposium on Information Theory (ISIT)*, pp. 1759–1763, 2009.
- [145] M. Torbatian, H. Najafi, and O. Damen, "Asynchronous Interference Channel: Degrees of Freedom and Interference Alignment," *CoRR*, vol. abs/1101.0275, [Online]. Available: <http://arxiv.org/abs/1101.0275>, 2011.
- [146] R. Tresch, M. Guillaud, and E. Riegler, "On the achievability of interference alignment in the K-user constant MIMO interference channel," in *IEEE/SP 15th Workshop on Statistical Signal Processing, 2009. SSP'09.*, pp. 277–280, IEEE, 2009.
- [147] D. Tse, P. Viswanath, and L. Zheng, "Diversity-multiplexing tradeoff in multiple-access channels," *IEEE Transactions on Information Theory*, vol. 50, no. 9, pp. 1859–1874, Sept 2004.
- [148] C. S. Vaze and M. K. Varanasi, "The Degrees of Freedom Regions of MIMO Broadcast, Interference, and Cognitive Radio Channels with No CSIT," *CoRR*, vol. abs/0909.5424, 2009.
- [149] C. S. Vaze and M. K. Varanasi, "The Degrees of Freedom Region and Interference Alignment for the MIMO Interference Channel with Delayed CSI," *CoRR*, vol. abs/1101.5809, [Online]. Available: <http://arxiv.org/abs/1101.5809>, 2011.
- [150] C. S. Vaze and M. K. Varanasi, "The Degrees of Freedom Region of the Two-User MIMO Broadcast Channel with Delayed CSI," *CoRR*, vol. abs/1101.0306, [Online]. Available: <http://arxiv.org/abs/1101.0306>, 2011.
- [151] S. Vishwanath, N. Jindal, and A. Goldsmith, "Duality, Achievable Rates, and Sum-Rate Capacity of MIMO Broadcast Channels," *IEEE Transactions on Information Theory*, pp. 2895–2909, October 2003.

- [152] P. Viswanath and D. Tse, "Sum capacity of the vector Gaussian broadcast channel and uplink-downlink duality," *IEEE Transactions on Information Theory*, pp. 1912–1921, August 2003.
- [153] C. Wang, T. Gou, and S. Jafar, "Interference Alignment through Staggered Antenna Switching for MIMO BC with no CSIT," in *Asilomar Conference on Signals, Systems and Computers*, November 2010.
- [154] C. Wang, T. Gou, and S. Jafar, "Multiple Unicast Capacity of 2-Source 2-Sink Networks," *CoRR*, vol. abs/1104.0954, [Online]. Available: <http://arxiv.org/abs/1104.0954>, 2011.
- [155] C. Wang, S. Jafar, S. Shamai, and M. Wigger, "Interference, Cooperation and Connectivity - A Degrees of Freedom Perspective," *CoRR*, vol. abs/1103.6060, [Online]. Available: <http://arxiv.org/abs/1103.6060>, 2011.
- [156] C. Wang, H. Papadopoulos, S. Ramprasad, and G. Caire, "Design and Operation of Blind Interference Alignment in Cellular and Cluster-based Systems," in *UCSD Information Theory and its Applications (ITA) Workshop 2011*, February 2011.
- [157] C. Wang, H. Papadopoulos, S. Ramprasad, and G. Caire, "Improved Blind Interference Alignment in a Cellular Environment using Power Allocation and Cell-Based Clusters," in *IEEE International Conference on Communications (ICC) 2011*, June 2011.
- [158] C. Wang and N. Shroff, "Beyond the butterfly—a graph-theoretic characterization of the feasibility of network coding with two simple unicast sessions," in *IEEE International Symposium on Information Theory, ISIT*, pp. 121–125, IEEE, 2007.
- [159] H. Weingarten, S. Shamai, and G. Kramer, "On the compound MIMO broadcast channel," in *Proceedings of Annual Information Theory and Applications Workshop UCSD*, January 2007.
- [160] H. Weingarten, Y. Steinberg, and S. Shamai, "The capacity region of the Gaussian MIMO broadcast channel," *IEEE Transactions on Information Theory*, vol. 52, pp. 3936–3964, September 2006.
- [161] Y. Wu, S. Shamai, and S. Verdú, "Degrees of Freedom of Interference Channel: a General Formula," in *Proceedings of International Symposium on Information Theory (ISIT) 2011*, 2011.
- [162] C. Yetis, T. Gou, S. Jafar, and A. Kayran, "On feasibility of interference alignment in MIMO interference networks," *IEEE Transactions on Signal Processing*, vol. 58, no. 9, pp. 4771–4782, 2010.
- [163] H. Yin, "Comments on Degrees of freedom region for K-user interference channel with M antennas," *CoRR*, vol. abs/1011.3812, [Online]. Available: <http://arxiv.org/abs/1011.3812>, 2010.
- [164] H. Yu, Y. Sung, H. Kim, and Y. Lee, "Adaptive beam tracking for interference alignment for multiuser time-varying MIMO interference channels," in *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pp. 3086–3089, IEEE, 2010.
- [165] W. Yu and J. Cioffi, "Sum capacity of Gaussian vector broadcast channels," *IEEE Transactions on Information Theory*, vol. 50, no. 9, pp. 1875–1892, Sept 2004.

- [166] Y. Zhu and D. Guo, "Isotropic MIMO interference channels without CSIT: The loss of degrees of freedom," *47th Annual Allerton Conference on Communication, Control, and Computing*, vol. abs/0910.2961, [Online]. Available: <http://arxiv.org/abs/0910.2961>, 2009.
- [167] Y. Zhu and D. Guo, "The degrees of freedom of MIMO interference channels without state information at transmitters," *CoRR*, vol. abs/1008.5196, [Online]. Available: <http://arxiv.org/abs/1008.5196>, 2010.