

Network Coding for Three Unicast Sessions: Interference Alignment Approaches

Abinesh Ramakrishnan[†], Abhik Das[‡], Hamed Maleki[†],
 Athina Markopoulou[†], Syed Jafar[†], Sriram Vishwanath[‡]
[†]EECS Dept., University of California, Irvine
[‡] ECE Dept., University of Texas, Austin

Abstract—We propose interference alignment techniques, originally developed for wireless interference channels, for the problem of network coding across unicast sessions. We describe two general approaches (namely, coding at the edge or in the middle of the network) and one specific example of each approach (namely, symbol extension method and ergodic alignment, respectively). We discuss the conditions for feasibility of alignment and their relation to network structure. We also compare alignment to alternative approaches. For three unicast sessions with min-cut one, we show that whenever alignment is possible, alternative approaches can also achieve half the min-cut. However, for more than three sessions and/or for min-cut per session greater than one, we show examples where alignment is necessary.

I. INTRODUCTION

The concept of interference alignment (IA) was originally developed for wireless interference channels [1], [2], [3]. The canonical example is a communication scenario where, regardless of the number of interferers, every user is able to access one half of the available spectrum free from the interference from other users. The key to interference alignment is the realization that alignment of signal spaces (in time, frequency, space and codes) is relative to the observer (receiver). Since every receiver sees a different picture, the signals may be constructed to cast overlapping shadows at the receivers where they constitute interference while they remain distinguishable at the receivers where they are desired.

In this work, we apply interference alignment techniques to the problem of network coding across different unicast sessions over directed acyclic graphs. This is a well-known open problem, for which optimal coding is not known and suboptimal constructive approaches are typically used. We refer to our approach, which combines network coding (NC) and alignment, as *network alignment (NA)*.

Our approach builds on the analogy depicted in Fig.1. Consider several unicast sessions over a wireline network; this can be thought of as equivalent to supporting the same unicast sessions over a wireless interference channel with the same linear transfer function. Network coding across sessions at intermediate nodes emulates superposition in the wireless channel. Essentially, the entire graph can be viewed as a channel, albeit a channel that is not given by nature, as it is the case in wireless, but determined by our routing and coding decisions. This has the advantage that it allows us to control the channel. However, it also has the disadvantage that it introduces spatial-correlation between end-to-end paths that




Fig. 1. Analogy between a graph with network coding and a wireless interference channel. Three unicast sessions are established on top of either network. Both systems can be represented by a linear transfer function and are amenable to alignment techniques.

share links among themselves. This correlation is not present in wireless channels, with high probability.

Inter-session network coding across unicast sessions introduces interference that prevents receivers from decoding the packets of the session they are really interested in. Interference alignment techniques can then be applied to guarantee rate of (at least) half the min-cut for each session, for any number of sessions transmitted over that network, subject to feasibility conditions. However, traditional IA techniques cannot be directly applied due to the correlation between end-to-end paths.

In this work, we are interested in developing systematic alignment approaches and algorithms for network coding across multiple unicasts. We are also interested in understanding their feasibility and performance depending on the network structure. We focus on a useful special case: network coding for three unicast sessions, which is the smallest, yet non-trivial, instance of the problem and which can be used as a building block for network coding across multiple unicasts.

The structure of the paper is as follows. Section II formulates the problem. Section III presents two general approaches (coding at the edge vs. coding in the middle) for network alignment and discusses their feasibility conditions. Section IV studies the throughput performance of network alignment compared to alternative approaches. Section V summarizes related work. Section VI concludes the paper.

II. FORMULATION OF THE THREE-UNICAST PROBLEM

Consider a network represented by a directed acyclic graph (DAG) $G = (V, E)$, where V is the set of nodes and E is the set of directed links. We assume that every directed link between a pair of nodes represents an error-free channel, and that the transmissions across different links do not interfere

with each other in any way. There are 3 source nodes, S_1, S_2, S_3 , and 3 destination nodes, D_1, D_2, D_3 and each S_i communicates only with D_i . The messages transmitted by different sources are assumed to be independent of each other. These messages are encoded and transmitted in form of symbols from finite field \mathbb{F}_p (p is a prime/prime power). For the sake of simplicity, we assume that every link in E has a capacity of one symbol (from \mathbb{F}_p) per channel use.

The coefficients for linear combination of symbols at each node come from \mathbb{F}_p . We consider these coefficients to be variables, say $\{\xi_1, \xi_2, \dots, \xi_s\}$ (s is a parameter dependent on the network topology), and define the vector $\underline{\xi} \triangleq [\xi_1 \ \xi_2 \ \dots \ \xi_s]$. A network coding scheme refers to choosing a suitable assignment for $\underline{\xi}$, from \mathbb{F}_p^s . Let the channel uses be indexed as $t = 1, 2, \dots$, and c_i be the min-cut for (S_i, D_i) . Then

$$\mathbf{y}_i(t) = \sum_{j=1}^3 \mathbf{M}_{ij}(\underline{\xi}) \mathbf{x}_j(t), \quad i = 1, 2, 3, \quad (1)$$

where $\mathbf{x}_i(t) \in \mathbb{F}_p^{c_i \times 1}$ is the input vector at S_i during the t^{th} channel use, $\mathbf{y}_i(t)$ is the $c_i \times 1$ output vector at D_i during the t^{th} channel use and $\mathbf{M}_{ij}(\underline{\xi})$ is the $c_i \times c_j$ transfer matrix between S_j and D_i . Note that the entries of $\mathbf{y}_i(t)$ and $\mathbf{M}_{ij}(\underline{\xi})$ are multivariate polynomials from the polynomial ring $\mathbb{F}_p[\underline{\xi}]$ for all i, j . Since D_i needs to decode only $\mathbf{x}_i(t)$ from $\mathbf{y}_i(t)$, the presence of transfer matrices $\mathbf{M}_{ij}(\underline{\xi})$, $i \neq j$, hinders the decodability (act as ‘‘interference’’) at every destination. We refer to these as ‘‘interference transfer matrices’’.

By the Max-flow-min-cut Theorem, S_i can transmit at most c_i symbols to D_i per channel use (here channel use refers to usage of one assignment of $\underline{\xi}$ from \mathbb{F}_p^s). The generalized Max-flow-min-cut Theorem, studied in [4], states that multiple unicast connections in G can achieve a maximum throughput of c_i for every source-destination pair (S_i, D_i) , iff there exists an assignment of $\underline{\xi}$ in \mathbb{F}_p^s , say $\underline{\xi}_0$, such that $\mathbf{M}_{ij}(\underline{\xi}_0) = 0$ for $i \neq j$ and $\mathbf{M}_{ii}(\underline{\xi}_0)$ is a full-rank matrix. However, there exists a broad class of networks for which such an assignment of $\underline{\xi}$ does not exist, thereby making multiple unicast at maximum throughput infeasible. Next we present our result that uses NA in this setting and also finds the condition for rate of half-the-min-cut to be achievable for each unicast session.

Special case: each of the three sessions has min-cut 1. The three input-output relations in (1) can be rewritten as:

$$\begin{aligned} y_1(t) &= m_{11}(\underline{\xi})x_1(t) + m_{12}(\underline{\xi})x_2(t) + m_{13}(\underline{\xi})x_3(t), \\ y_2(t) &= m_{21}(\underline{\xi})x_1(t) + m_{22}(\underline{\xi})x_2(t) + m_{23}(\underline{\xi})x_3(t), \\ y_3(t) &= m_{31}(\underline{\xi})x_1(t) + m_{32}(\underline{\xi})x_2(t) + m_{33}(\underline{\xi})x_3(t), \end{aligned}$$

where $x_i(t), y_i(t)$ and $m_{ij}(\underline{\xi})$ are the scalar equivalents of $\mathbf{x}_i(t), \mathbf{y}_i(t)$ and $\mathbf{M}_{ij}(\underline{\xi})$ respectively. Moreover, we have $x_i(t) \in \mathbb{F}_p$ and $y_i(t), m_{ij}(\underline{\xi}) \in \mathbb{F}_p[\underline{\xi}]$ for $i, j = 1, 2, 3$. Note that $m_{ii}(\underline{\xi})$, $i = 1, 2, 3$, are non-trivial polynomials. Also by construction, $m_{ii}(\underline{\xi})$ cannot be a non-zero constant and $m_{ii}(\underline{\xi}) \not\equiv c m_{ij}(\underline{\xi})$, where $c \in \mathbb{F}_p$. Thus, these non-trivial polynomials are exclusive functions of $\underline{\xi}$. We refer to $m_{ij}(\underline{\xi})$, $i, j = 1, 2, 3$, as ‘‘network transfer functions’’. We also




Fig. 2. **Network alignment approaches.** **Left: alignment by coding in the middle of the network.** Here, intermediate nodes perform linear network coding so as to guarantee alignment at the receivers. One construction is the ergodic scheme that uses two time-slots, described in this paper and inspired by [5]. **Right: Alignment by coding at the edge of the network.** Here, intermediate nodes perform random linear network coding, and the sources are responsible for coding their symbols so as to guarantee alignment. We proposed one construction in [7], which follows the asymptotic scheme of [1] and enables all sessions to achieve asymptotically half their min-cut.

refer to polynomials $m_{ii}(\underline{\xi})$, $i = 1, 2, 3$, as ‘‘direct paths’’ and polynomials $m_{ij}(\underline{\xi})$, $i \neq j$, as ‘‘interference paths’’.

III. NETWORK ALIGNMENT APPROACHES

A. Lessons learned from the wireless interference alignment

Various approaches have been pursued to achieve interference alignment in wireless networks. In particular, within the linear framework, two distinct approaches have been used extensively. The first is the asymptotic IA (also known as ‘‘symbol-extension’’) approach, originally introduced in [1], where IA is achieved in the limit of large number of signaling dimensions. The second approach known as ergodic IA, first introduced in [5], has an opportunistic flavor. While the asymptotic approach is known to be much more widely applicable, the ergodic IA approach is also applicable to a fairly broad class of commonly studied wireless networks and is in general much more efficient – in many cases achieving theoretical capacity region as shown in [6].

Translated to the network coding context, the tradeoffs between the two approaches show interesting new aspects. Fig. 2 depicts at high-level the two general approaches for network alignment: coding at the edge of the network (e.g., using the asymptotic scheme in [7], [1]) or coding in the middle of the network (e.g., using an ergodic scheme similar to [5]). We discuss both of these approaches in more detail.

B. Coding at the edge of the network

This approach is closer to the wireless paradigm. It uses only random linear network coding at the intermediate nodes and essentially creates a random linear channel between the sources and destination nodes. The intelligence and the processing burden is pushed to the edge, much like the wireless setting where the channel is decided by nature and all the

coding/signal processing algorithms are applied only at the source and destination nodes. In addition to the sophisticated processing required at the end nodes, this approach often requires either a very large alphabet size or very large amounts of data due to its asymptotic character (in channel uses). Nevertheless, it is a very powerful approach at least for establishing the theoretical performance limits.

In recent work [7], we applied the symbol extension scheme [1] (over finite fields as opposed to reals). The scheme achieves asymptotic alignment *i.e.*, for large number of symbols, by pre-coding only at the sources of a network, while the rest of the network performs random linear network coding. The pre-coding matrices depend on the network transfer functions $m_{ij}(\underline{\xi})$'s; the reader is referred to [7] for the details on their construction. It is worth noting that the dependencies between $m_{ij}(\underline{\xi})$'s due to the network structure impose some conditions for symbol-extension alignment to be possible. This is unlike the wireless setting, where $m_{ij}(\underline{\xi})$'s are channel gains and typically assumed to be independent. Let us define:

$$\begin{aligned} a(\underline{\xi}) &= m_{12}(\underline{\xi})m_{23}(\underline{\xi})m_{31}(\underline{\xi}), \\ b(\underline{\xi}) &= m_{21}(\underline{\xi})m_{13}(\underline{\xi})m_{32}(\underline{\xi}). \end{aligned}$$

In [7], we derived sufficient conditions for the feasibility of asymptotic alignment over networks with network coding and three unicast sessions. Assuming the polynomials $m_{ij}(\underline{\xi})$'s to be non-trivial, these state that $\forall n$ and $\forall p_i, q_j \in \mathbb{F}_p$, $i, j = 0, 1, \dots, n$, the following conditions should hold:

$$m_{11}(\underline{\xi}) \neq \frac{m_{12}(\underline{\xi})m_{31}(\underline{\xi}) \sum_{i=0}^n p_i (a(\underline{\xi})/b(\underline{\xi}))^i}{m_{32}(\underline{\xi}) \sum_{j=0}^n q_j (a(\underline{\xi})/b(\underline{\xi}))^j}, \quad (2)$$

$$m_{22}(\underline{\xi}) \neq \frac{m_{21}(\underline{\xi})m_{32}(\underline{\xi}) \sum_{i=0}^n p_i (a(\underline{\xi})/b(\underline{\xi}))^i}{m_{31}(\underline{\xi}) \sum_{j=0}^n q_j (a(\underline{\xi})/b(\underline{\xi}))^j}, \quad (3)$$

$$m_{33}(\underline{\xi}) \neq \frac{m_{23}(\underline{\xi})m_{31}(\underline{\xi}) \sum_{i=0}^n p_i (a(\underline{\xi})/b(\underline{\xi}))^i}{m_{21}(\underline{\xi}) \sum_{j=0}^n q_j (a(\underline{\xi})/b(\underline{\xi}))^j}. \quad (4)$$

In order to get a better intuition about the relation between these conditions and the network structure, let us consider a subset of these conditions that is obtained by considering some special, small values of the constants p_i, q_j . In particular, the conditions at the left are obtained for $p_0 = 1, q_0 = 1$, and the conditions at the right are obtained for $p_0 = 1, q_1 = 1$ or $p_1 = 1, q_0 = 1$; the rest p_i 's, q_j 's are set to zero:

$$m_{11}(\underline{\xi}) \neq \frac{m_{12}(\underline{\xi})m_{31}(\underline{\xi})}{m_{32}(\underline{\xi})}, \quad m_{11}(\underline{\xi}) \neq \frac{m_{21}(\underline{\xi})m_{13}(\underline{\xi})}{m_{23}(\underline{\xi})}, \quad (5)$$

$$m_{22}(\underline{\xi}) \neq \frac{m_{21}(\underline{\xi})m_{32}(\underline{\xi})}{m_{31}(\underline{\xi})}, \quad m_{22}(\underline{\xi}) \neq \frac{m_{12}(\underline{\xi})m_{23}(\underline{\xi})}{m_{13}(\underline{\xi})}, \quad (6)$$

$$m_{33}(\underline{\xi}) \neq \frac{m_{23}(\underline{\xi})m_{31}(\underline{\xi})}{m_{21}(\underline{\xi})}, \quad m_{33}(\underline{\xi}) \neq \frac{m_{32}(\underline{\xi})m_{13}(\underline{\xi})}{m_{12}(\underline{\xi})}. \quad (7)$$

The above six conditions require that the co-factor of any off-diagonal term of the 3×3 network transfer matrix $[m_{ij}(\underline{\xi})]$ is not the zero polynomial. Let us consider one such off-diagonal term, *e.g.*, $m_{13}(\underline{\xi})$ with its cofactor:

$$\det \begin{bmatrix} m_{21}(\underline{\xi}) & m_{22}(\underline{\xi}) \\ m_{31}(\underline{\xi}) & m_{32}(\underline{\xi}) \end{bmatrix} = m_{21}(\underline{\xi})m_{32}(\underline{\xi}) - m_{22}(\underline{\xi})m_{31}(\underline{\xi}).$$

The condition $m_{21}(\underline{\xi})m_{32}(\underline{\xi}) - m_{22}(\underline{\xi})m_{31}(\underline{\xi}) \neq 0$ essentially states that aligning interference at one receiver (in this case, interference $m_{31}(\underline{\xi}), m_{32}(\underline{\xi})$ from S_1, S_2 at D_3) should not have the undesired side-effect of aligning the signal with the interference at other receivers (in this case, interference $m_{21}(\underline{\xi})$ and signal $m_{22}(\underline{\xi})$ at D_2). In terms of network structure, the conditions states that not all paths for the four commodities involved (from S_1 to D_2, S_2 to D_2, S_1 to D_3 , and S_2 to D_3) should go through the same bottleneck edge. In other words, the 2×2 structure should have rank 2.

Interestingly, the conditions about the non-zero off-diagonal co-factors coincide with the necessary conditions for achieving a rate more than $1/3$ per user through any achievable scheme, as shown previously in [8] based on Shannon theoretic arguments for the 3-user wireless interference network setting. It is also important to note that although these conditions are always met (with high probability) in wireless channels, they may not hold for all DAGs. In general, not all networks can be aligned. On the positive side, there are graphs over which aligned sessions can achieve more than the minimum guarantee (rate of half-the-min-cut). Therefore, an open research question is to characterize the feasibility and performance of alignment and their relation to network structure.

C. Coding in the middle of the network

This approach is depicted on the left side of Fig. 2 and can include several algorithms. Here we outline one such algorithm which we call *ergodic network alignment* because it is inspired by the ergodic alignment in wireless [5]. Consider again the special case of 3 unicasts with min-cut one each, and choose the coding coefficients in two time slots so that they lead to the following two network transfer matrices:

$$M^{(t)} = \begin{bmatrix} m_{11}^{(t)} & m_{12} & m_{13} \\ m_{21} & m_{22}^{(t)} & m_{23} \\ m_{31} & m_{22} & m_{33}^{(t)} \end{bmatrix}, \quad t = 1, 2.$$

Essentially we have picked the coding coefficients so that the interference polynomials m_{ij} , $i \neq j$, remain the same in both time slots, while m_{ii} 's change values across time slots. If we transmit the same symbol in both time slots, we can subtract the interference terms and decode that one symbol, thus achieving rate of half the min-cut per user. This alignment scheme is feasible if each m_{ii} is not a function (*any* function) of m_{ij} , $i \neq j$. Notice that this is a much stronger requirement than the conditions for the feasibility of the asymptotic scheme, which requires only that m_{ii} 's are not

Network	# nodes	# edges
Telstra (AUS)	108	153
Sprintlink (USA)	315	972
EBONE (EU)	87	161
Tiscali (EU)	161	328
Exodus (USA)	79	147
Abovenet (USA)	141	374
ABILENE	11	14
BELNET	15	27
GEANT	23	37

TABLE I
REAL TOPOLOGIES OF INTERNET SERVICE PROVIDERS (ISPs), AS
MEASURED BY THE ROCKETFUEL PROJECT.


specific functions of m_{ij} , $i \neq j$ (Eq.(2)-(4)).

This approach is particularly attractive in the network coding setting because it operates in two time slots and thus can lead to practical solutions in terms of the required field size and number of symbols. Moreover, unlike the wireless setting where the ergodic IA approach entails inordinate delays waiting for the optimal channel conditions to facilitate simple IA, in the network coding setting the channel conditions are determined by the operations at the intermediate nodes and hence can be controlled. This simplicity comes at the cost of introducing intelligence in the network, and depending on the size of the network this optimization may be difficult.

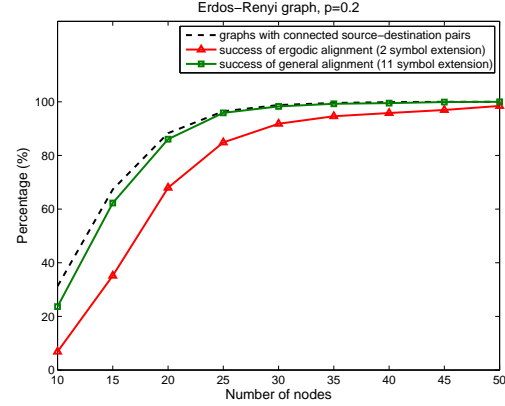
D. Simulation Results

In order to understand how “mild” the feasibility conditions (for alignment) are in various topologies, we perform simulations. We showcase some representative results in Fig. 3. First, we simulated Erdos-Renyi graphs, for varying number of nodes n and probability of edge p ; the results are reported in Fig. 3(a)-(b). We also considered nine real network topologies of well-known ISPs, measured by Rocketfuel and summarized in Table I; the results are shown in Fig. 3(c).


For each topology, we randomly pick three pairs of source-destination pairs, (S_i, D_i) , $i = 1, 2, 3$. Because we assume that they have min-cut of 1 each, we fix a unique path between every source-destination pair (S_i, D_j) , and we perform coding on the resulting directed subgraph. There are several ways we can fix a path; we choose the shortest path (if there are multiple shortest paths, we arbitrarily choose one of the paths). The presence of paths (S_i, D_j) , $i \neq j$, acts as “interference”, and therefore alignment needs to be done at the destinations. We consider an 11-symbol extension for the asymptotic alignment scheme; we randomly choose the coding coefficients for 11 channel uses, construct the precoding matrices and check if the alignment is possible using these matrices. We also consider a 2-symbol extension ergodic alignment scheme, where we examine if each of the diagonal transfer functions $m_{ii}(\xi)$ has a coding variable, not present in the transfer functions $m_{ij}(\xi)$, $i \neq j$. Note that this choice is overly restrictive, since ergodic alignment only requires that $m_{ii}(\xi)$ is not a function of $m_{ij}(\xi)$, $i \neq j$, which is already more restrictive



(a) Random graphs, $n = 15$, varying p



(b) Random graphs, $p = 0.2$, varying n



(c) Real Internet topologies

Fig. 3. Testing the feasibility of alignment in various topologies.

than the conditions for asymptotic alignment (which requires that $m_{ii}(\xi)$ is not a specific function of $m_{ij}(\xi)$, $i \neq j$). We weigh the performance of the alignment schemes alongside the percentage of graphs with connected (S_i, D_i) pairs.

One can make the following observations based on Fig.3. First, alignment is feasible in large percentage of topologies. In

random graphs, the symbol-extension scheme succeeds almost whenever the source-destination pairs are connected. In 7 out of 9 real topologies, the asymptotic alignment scheme is feasible close to 100% of the times. For the remaining two topologies, a closer inspection showed that these topologies had rank 1, which means that no approach can achieve half the min-cut anyway. Second, the 2-symbol scheme is feasible in smaller percentage of the topologies, due to the simple scheme we considered. The percentage of success for the simple scheme we simulated is already high but provides only a lower bound to the performance of the ergodic approach. Third, whenever alignment failed, we tested whether the six “small” conditions (Eq.(5)-(7)) are violated. This turned out to always be the case, as reported in Fig.3(c): see that the percentage of times that the 11-symbol extension alignment fails is the same with the percentage of times that the “small” conditions are violated. Based on this observation and from the intuitive interpretation of the “small” conditions, we conjecture that the six “small” conditions are sufficient for the feasibility of alignment in DAGs with three unicast sessions.


IV. COMPARING ALIGNMENT TO ALTERNATIVES

A. Three unicast sessions, each with $\text{min-cut}=1$


Let us first discuss some illustrative examples. Consider the canonical example of the extended butterfly, shown in Fig. 4(a). In this case, routing can achieve rate of $1/3$ since all flows go through the bottleneck link. Network coding achieves rate of 1 per flow if all side links (defined as links between S_i and R_j , $i \neq j$) are available. Let us consider the same example (all flows still go through a bottleneck link) but now only a subset of the side links is present. If only one receiver has side links, then two of the receivers have the same view, and alignment becomes impossible. If two or three receivers have side links, then depending on which links those are, network alignment may be possible. The question we are interested in is: can other approaches also achieve rate of half?

Fig. 4(b) illustrates a case where two receivers have a side link each, but alignment is not possible; e.g., one can check that the “small” conditions (Eq.(5)-(7)) are violated. Fig. 4(c) illustrates another case where two receivers have a side link each. Notice that compared to the previous example, receivers 2 and 3 have switched place. Alignment is possible in this case and can achieve half the min-cut per session. However, the side links now form a butterfly substructure (between session 1 and 3). Therefore, it is possible to achieve the same rate by time-sharing between the butterfly and session 2.


Fig. 4(d) illustrates a case where all three receivers have a side link and there is no butterfly (for any 2 sessions) present in the network. Alignment can achieve half the min-cut, which is optimal in this case. However, it is also possible to achieve half the min-cut per session by carefully choosing the coefficients at the nodes $1'$, $2'$ and $3'$, so as to cancel out one component of the interference for each session. For example, node $1'$ can pick $q = -cp$ so as to cancel z and allow receiver 1 to only see equations in two unknowns (x, y), thus making it possible to solve for its own message x over two time slots.




(a) The extended butterfly: routing achieves $1/3$, alignment achieves $1/2$, and network coding achieves rate 1 per session.



(b) Alignment is not feasible



(c) Alignment is feasible but a butterfly is also present



(d) Alignment is feasible but coding without alignment can also achieve half the min-cut

Fig. 4. Examples for 3 unicast sessions, min-cut 1 where all 3 sessions go through a single bottleneck. In all these examples, whenever alignment is possible, an alternative approach can also achieve $1/2$ the rate.

The above intuition generalizes to a more general statement.

Theorem IV.1. *Consider a DAG with three unicast sessions, each with min-cut of 1. Whenever network alignment can achieve rate of $1/2$ per session, there exists an alternative approach that can also achieve rate of $1/2$ per session.*

The alternative approaches include: routing, packing butterflies, random linear network coding, or other network coding strategies that do not require the alignment strategy.¹

¹Notice that alignment may be achieved by coding in the middle or at the edge of the network. What distinguishes “alignment” from “non-alignment” approaches is that alignment provides the receivers with lesser number of equations than unknowns, but with aligned interference.

Proof: The common routing rate r that can be guaranteed per flow is upper bounded by the sparsity bound S . We consider two cases: $S \leq 1/2$ and $S > 1/2$. Notice that the sparsity bound S in this setup can only take one of the following values $\{1/3, 1/2, 2/3, 1\}$, provided that each session demands a rate equal to its corresponding min-cut ($= 1$).

Case I: Let us consider networks where $S = 1/3 < 1/2$. In these networks, we have $r \leq S = 1/3$. Therefore, routing cannot achieve half the min-cut and alignment needs only be compared against coding alternatives. $S = 1/3$ in this setup means that all paths of the three commodities traverse the same bottleneck edge. The intuition of the examples we discussed earlier still applies with the difference that links should be interpreted as paths from/to the sources/receivers.

In general, a network transfer function $m_{ij}(\xi)$ is just a sum of several monomials, where each monomial represents a path from source j to sink i . It is possible to express $m_{ij}(\xi) = d_{ij}(\xi) + c_{ij}(\xi)$, where $d_{ij}(\xi)$ is the polynomial resulting from the sub graph G' formed by removing any edge that doesn't belong to a *direct* path connecting source j to its corresponding sink i ($\forall j = 1, 2, 3$) in G ; $c_{ij}(\xi)$ is the polynomial resulting from all the paths that traverse through those edges that were neglected in G' ; we refer to the corresponding paths as *cross* or *side* paths. Note that $d_{ii}(\xi) \neq 0$ and $c_{ii}(\xi) = 0$, $i = 1, 2, 3$. Consider the matrices $D(\xi) = [d_{ij}(\xi)]$ and $C(\xi) = [c_{ij}(\xi)]$, and note that $M(\xi) = [m_{ij}(\xi)] = D(\xi) + C(\xi)$.

Because of the single bottleneck link between all direct paths in this case, $\text{rank}[D(\xi)] = 1$ and the rate achievable by routing is $1/3$ per session. Also, note that the routing rate doesn't depend on any of the $c_{ij}(\xi)$'s but the side paths $c_{ij}(\xi)$'s do determine the feasibility of alignment. It is quite obvious that if more than one row of $C(\xi)$ is all zeros, then two or more sinks will be seeing the same perspective and it would be impossible to achieve a rate of more than $1/3$ per session. If one of the rows (say row i') alone is made of all zero elements, then we need to have $c_{i'j}(\xi) \neq 0$ and $c_{i'k}(\xi) \neq 0$ in order to satisfy the feasibility conditions (Eq.(5)-(7))

$$m_{jj}(\xi) \neq \frac{m_{kj}(\xi) m_{ji'}(\xi)}{m_{i'k}(\xi)} \quad \forall j \neq k \neq i',$$

and this results in forming a butterfly structure between sessions j and k . An example was shown in Fig. 4(c): alignment is feasible but there is also a butterfly structure. In Fig. 4(b), the side paths do not form a butterfly but they do not satisfy the alignment feasibility conditions either: *i.e.*, $m_{11}(\xi) \equiv (m_{12}(\xi)m_{31}(\xi))/m_{32}(\xi)$ and $m_{22}(\xi) \equiv (m_{21}(\xi)m_{32}(\xi))/m_{31}(\xi)$. If $C(\xi)$ has a non-zero entry in every row, as in Fig.4(d), it is possible to achieve half the min-cut, without alignment, by carefully choosing the coefficients at the node where paths corresponding to the non zero $c_{ij}(\xi)$'s join with the path corresponding to $d_{ij}(\xi)$'s.

Case II: Let us consider networks where $S \geq 1/2$. If routing can achieve $r \geq 1/2$, then we are done. The cases that require investigation are the ones where $r < 1/2 \leq S$; in these cases, we will describe a scheme that uses network coding (without

alignment) at carefully selected nodes in the middle of the network, and guarantees rate of $1/2$ for every session. In fact, it achieves this goal by using only the direct paths.

Let $\mathcal{P}_i = \{P_i^{(n)} : n = 1, 2, \dots, N_i\}$ represent the set of all direct paths from source i to receiver i . Choose a path for each session, say $P_1^{(\alpha)} \in \mathcal{P}_1$, $P_2^{(\beta)} \in \mathcal{P}_2$ and $P_3^{(\gamma)} \in \mathcal{P}_3$ where $\alpha \in \{1, \dots, N_1\}$, $\beta \in \{1, \dots, N_2\}$ and $\gamma \in \{1, \dots, N_3\}$. We are in the case where the rate achievable by routing is smaller than $1/2$; this means that the three paths $P_1^{(\alpha)}$, $P_2^{(\beta)}$, $P_3^{(\gamma)}$ traverse through a bottleneck edge. But we are also in the case where $S > 1/3$; this ensures that there is another path $P' \in \mathcal{P}_1$ or \mathcal{P}_2 or \mathcal{P}_3 such that there is no bottleneck edge through which all four of the paths ($P_1^{(\alpha)}$, $P_2^{(\beta)}$, $P_3^{(\gamma)}$ and P') will traverse. Without loss of generality we can assume that $P' \in \mathcal{P}_1$. P' , $P_2^{(\beta)}$, $P_3^{(\gamma)}$ should have a bottleneck edge so that routing can't achieve half rate. Any edge that doesn't belong to any of the previously defined four paths can be ignored for the design of the scheme we are about to describe.

The scheme involves sharing the network between the sessions over time. In the first time slot, the network is used by a pair of sessions; during the second time slot the session that was left out in the first time slot gets the network resources to itself. To achieve rate of half the min-cut, we need to make sure that the two sessions, sharing the network in the first time slot, are able to decode their respective messages within that time slot. This is done by carefully choosing the coding coefficients at some nodes. To pick the pair of sessions that would be active in the first time slot we need to look at the first and last edge where all direct paths of any two sessions overlap. Such an overlap is always present because of the bottlenecks described earlier. Let edge e_f be the first edge, where all paths of sessions (i, j) overlap, and edge e_l be the final edge, where sessions (j, k) overlap at the end, then we will choose session (i, k) to be the active pair in the first time slot.

If session 1 is one of the active sessions in the first time slot, then we know that there are two paths $P_1^{(\alpha)}$ and P' joining at some node after e_l . We can choose the coefficients at the tail node of the edge e_l such that the message component of session 1 is eliminated and only the component of the other session is present at the head of the edge e_l . Thus the sink of the other session gets its message without any interference and session 1 can still decode its message by using the component of the other session, from the head of edge e_l , as an antidote at the edge where the two paths of session 1 meet. If session 1 is not among the pair of active sessions in the first time slot, then we can use the same idea to cancel out the interference before delivering the messages to the respective sinks of the sessions. But here, we choose the coefficients at the tail node of edge e_l to cancel out the interference component of session i and the coefficients at the tail node of the edge right before e_l where the paths of session i and session k overlap. This ensures that each session that is active in the first time slot gets its corresponding message delivered within that time slot. This scheme succeeds in achieving rate of half-the-min-cut per session without using any alignment techniques. ■

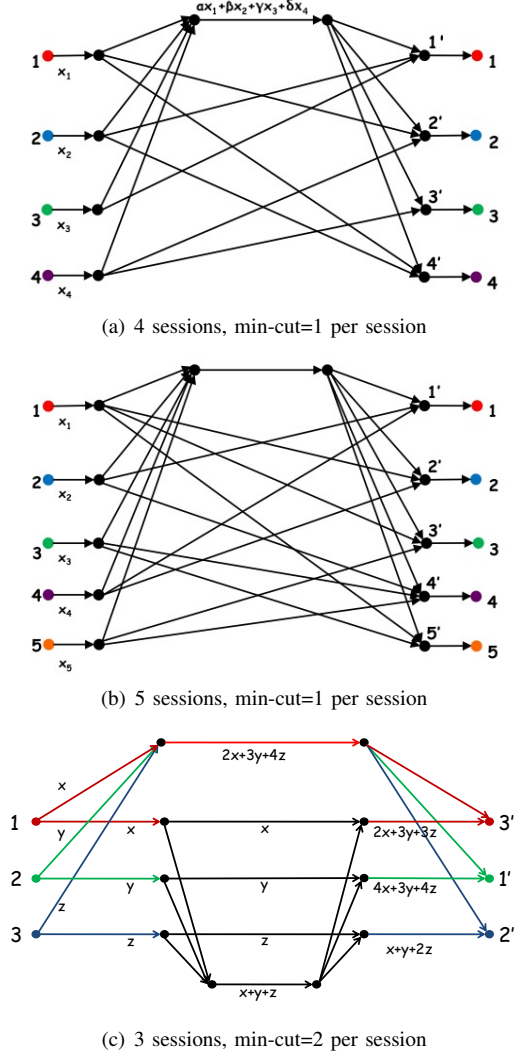


Fig. 5. Examples that require alignment in order to achieve half the min-cut per session. They have either more than three unicast sessions or min-cut greater than one.

B. More than three sessions or min-cut greater than 1

Fig. 5(a)-(c) show examples where the number of sessions is $K > 3$ or the min-cut is greater than one. In all these cases alignment is required in the following sense: the maximum rate is $1/2$ the min-cut, alignment achieves it and no other method (packing butterflies or coding without alignment) can achieve it. Indeed, in Fig. 5(a), receiver 3 has 2 equations with 4 unknowns over 2 time slots. (Notice that the side links convey the redundant information over 2 time slots.) Even if the sink nodes are considered to be present at $1', 2', 3'$ and $4'$, we can observe that over 2 time slots node $3'$ receives only 3 independent equations with 4 unknowns. So it needs an alignment technique to decode its message. This effect is amplified for 5 sessions example in Fig. 5(b), where the nodes $1', 2', 3'$ and $5'$ effectively get 4 equations in 5 unknowns over two time slots. In Fig. 5(c) it can be observed that the equations carried by the incoming edges of node $D_i, \forall i = 1, 2, 3$, needs

to have their interference component aligned in order for the receiver to be able to decode its corresponding message.

V. RELATED WORK

Inter-session network coding. Since the early work on network coding [9], [10], [4], it became clear that inter-session network coding, which includes the practical case of multiple unicasts, is a hard problem. A sufficient condition for optimality of linear inter-session network coding was developed in [4]. However, scalar or even vector linear network coding alone has been shown to be insufficient for optimal inter-session network coding [11]. Even computing the inner/outer bounds of the rate region cannot be computed in practice [12], [13]. In short, optimal inter-session coding is today an open research problem. Only heuristic/suboptimal approaches are considered in practice, often without even performance guarantees. A key difficulty is the combinatorial nature of the problem: one needs to enumerate and select from all possible subsets of the flows that can be combined together and from all possible locations in the graph where coding/decoding should be performed. To make the problem tractable, some heuristic approaches consider only very restrictive classes of codes, either in terms of coding operations allowed (*e.g.*, XOR operations) and/or in terms of coding subgraph (*e.g.*, all packets are coded/decoded hop-by-hop in [14]; butterfly structures consider two hops [15]; tiling of patterns is considered in [16]). The most restrictive case is XOR coding between pairs of flows: this allows for a network flow formulation of the problem [17] known as butterfly packing. A nice survey of state-of-the-art approaches and a novel evolutionary approach that selects among different suboptimal policies can be found in [18].

Interference alignment in wireless networks. The IA concept was originally developed for interference channels [1], [19]. IA schemes have been found for a variety of networks including X networks [2], [3], compound broadcast channel [20], [21], interference networks [1], [22], [23], [24], [25], [26], [5], [27], cellular networks [28], [29], multihop (relay) networks [30], bidirectional relay networks [31] and over wireless networks supported by a wired backbone [32]. These include linear alignment schemes over signal spaces (introduced in [3], [20]), signal level and lattice alignment schemes (introduced in [22], [25]), asymptotic alignment schemes over a large number of dimensions (introduced in [1]), propagation delay based alignment [33], [34], asymmetric complex signaling schemes (introduced in [29]), rational/irrational scaled lattice alignment schemes (introduced in [26], [27]), interference alignment and cancellation schemes (introduced in [32]) and ergodic alignment schemes (introduced in [5]). The ideas of interference alignment and cancellation have also been recently implemented in a practical system [32].

Network Alignment: from wireless channels to graphs. In our recent work [7], we introduced the use of alignment techniques for network coding across multiple unicasts, using the symbol extension method developed in [1]. In this paper, we revisited this work and we discussed in more depth the

feasibility conditions and the throughput benefit, especially for the case of three unicast sessions.

The role of alignment has also been recognized in the context of distributed storage, for solving the systematic repair bandwidth minimization problem [35], [36], [37]. The problem was solved through elegant and explicit alignment solutions for all (k, n) MDS codes such that $n \geq 2k$. The general solution for all (k, n) was found recently in [38], [35] and requires the use of the asymptotic alignment scheme of [1].

VI. CONCLUSION

We introduced the use of interference alignment techniques, originally developed for wireless interference channel, into the problem of network coding across three unicast sessions. We described two general approaches, *i.e.*, coding at the edge or in the middle of the network, and two specific examples of each approach, *i.e.*, the symbol extension method and ergodic alignment, respectively. We discussed the feasibility conditions and their relation to network structure and compared alignment to alternative approaches. For three unicast sessions with min-cut one, we show a negative result: whenever alignment is possible, alternative approaches can also achieve half the min-cut. However, for more than three sessions and/or for min-cut per session greater than one, we show examples where alignment is necessary. Directions for future work include: further investigation of the relation between network structure and feasibility/performance of alignment; and the design of practical network alignment algorithms.

REFERENCES

- [1] V. Cadambe and S. Jafar, "Interference alignment and the degrees of freedom of the k user interference channel," *IEEE Trans. on Information Theory*, vol. 54, no. 8, pp. 3425–3441, Aug. 2008.
- [2] M. Maddah-Ali, A. Motahari, and A. Khandani, "Communication over MIMO X channels: Interference alignment, decomposition, and performance analysis," in *IEEE Trans. on Information Theory*, 2008, pp. 3457–3470.
- [3] S. Jafar and S. Shamai, "Degrees of freedom region for the MIMO X channel," *IEEE Trans. on Information Theory*, vol. 54, no. 1, pp. 151–170, Jan. 2008.
- [4] R. Koetter and M. Medard, "An algebraic approach to network coding," *Networking, IEEE/ACM Transactions on*, vol. 11, no. 5, pp. 782–795, Oct. 2003.
- [5] B. Nazer, M. Gastpar, S. A. Jafar, and S. Vishwanath, "Ergodic interference alignment," June 2009.
- [6] S. A. Jafar, "The ergodic capacity of interference networks," e-print, arxiv:0902.0838.
- [7] A. Das, S. Vishwanath, S. Jafar, and A. Markopoulou, "Network coding for multiple unicasts: An interference alignment approach," *Proc. IEEE International Symposium on Information Theory. arXiv:1008.0235*, 2010.
- [8] V. Cadambe and S. Jafar, "Parallel gaussian interference channels are not always separable," *IEEE Trans. on Info. Theory*, vol. 55, pp. 3983–3990, Sep. 2009.
- [9] R. Ahlswede, N. Cai, S.-Y. Li, and R. Yeung, "Network information flow," *Information Theory, IEEE Transactions on*, vol. 46, no. 4, pp. 1204–1216, Jul 2000.
- [10] S.-Y. Li, R. Yeung, and N. Cai, "Linear network coding," *Information Theory, IEEE Transactions on*, vol. 49, no. 2, pp. 371–381, Feb 2003.
- [11] R. Dougherty, C. Freiling, and K. Zeger, "Insufficiency of linear coding in network information flow," *IEEE Transactions on Information Theory*, vol. 51, no. 8, pp. 2745 – 2759, Aug. 2005.
- [12] X. Yan, R. W. Yeung, and Z. Zhang, "The capacity region for multi-source multi-sink network coding," *Information Theory, IEEE International Symposium on*, pp. 116 –120, June 2007.
- [13] R. Yeung, S. Li, N. Cai, and Z. Zhang, "Network coding theory part ii: multiple source," *Foundations and Trends in Communications and Information Theory*, vol. 2, no. 5, pp. 330–381, 2005.
- [14] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft, "Xors in the air: practical wireless network coding," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 4, pp. 243–254, 2006.
- [15] S. Omiwade, R. Zheng, and C. Hua, "Butterflies in the mesh: lightweight localized wireless network coding," in *Proc. of NetCod*, Jan. 2008.
- [16] M. Effros, T. Ho, and K. S., "A tiling approach to network code design for wireless networks," in *Proc. of IEEE Information Theory Workshop (ITW)*, Mar. 2006.
- [17] D. Traskov, N. Ratnakar, D. Lun, R. Koetter, and M. Medard, "Network coding for multiple unicasts: An approach based on linear optimization," July 2006, pp. 1758–1762.
- [18] M. Kim, U.-M. O'Reilly, M. Medard, and D. Traskov, "An evolutionary approach to inter-session network coding," in *Proc. of IEEE INFOCOM*, Apr. 2009.
- [19] A. Motahari and A. Khandani, "Capacity bounds for the gaussian interference channel," *Information Theory, IEEE Transactions on*, vol. 55, no. 2, pp. 620–643, Feb. 2009.
- [20] H. Weingarten, S. Shamai, and G. Kramer, "On the compound MIMO broadcast channel," in *Proceedings of Annual Information Theory and Applications Workshop UCSD*, Jan 2007.
- [21] T. Gou, S. A. Jafar, and C. Wang, "On the degrees of freedom of finite state compound wireless networks - settling a conjecture by weingarten et. al," 2009, <http://arxiv.org/abs/0909.4177>.
- [22] G. Bresler, A. Parekh, and D. Tse, "Approximate capacity of the many-to-one interference channel," *Allerton Conference*, Sep. 2007.
- [23] V. R. Cadambe, S. A. Jafar, and S. Shamai, "Interference alignment on the deterministic channel and application to fully connected gaussian interference networks," *Information Theory, IEEE Transactions on*, vol. 55, no. 1, pp. 269–274, Jan. 2009.
- [24] T. Gou and S. A. Jafar, "Capacity of a class of symmetric SIMO Gaussian interference channels within $O(1)$," *arxiv.org*, vol. abs/0905.1745, 2009, <http://arxiv.org/abs/0905.1745>.
- [25] S. Sridharan, A. Jafarian, S. Vishwanath, and S. A. Jafar, "Capacity of symmetric K-user gaussian very strong interference channels," *Proceedings of IEEE Globecom*, pp. 1–5, Dec 2008.
- [26] R. Etkin and E. Ordentlich, "On the degrees-of-freedom of the K-user Gaussian interference channel," *arxiv.org*, vol. abs/0901.1695, 2009.
- [27] A. S. Motahari, S. O. Gharan, M. A. Maddah-Ali, and A. K. Khandani, "Forming pseudo-mimo by embedding infinite rational dimensions along a single real line: Removing barriers in achieving the dofs of single antenna systems," *arxiv.org*, vol. abs/0908.2282, 2009, <http://arxiv.org/abs/0908.2282>.
- [28] C. Suh and D. Tse, "Interference alignment for cellular networks," in *Proceedings of 40th Annual Allerton Conference on Communication, Control and Computing*, Sep 2008.
- [29] V. R. Cadambe, S. A. Jafar, and C. Wang, "Interference alignment with asymmetric complex signaling - settling the host-madsen-nosratinia conjecture," in *arXiv:0904.0274*, 2009.
- [30] S.-W. Jeon, S.-Y. Chung, and S. A. Jafar, "Degrees of freedom of multi-source relay networks," 2009.
- [31] N. Lee and J. Lim, "A novel signaling for communication for MIMO Y channel: Signal space alignment for network coding," in *Proc. Int. Symp. Information Theory*, June 2009.
- [32] S. Gollakotta, S. Perli, and D. Katabi, "Interference alignment and cancellation," *Proceedings of ACM SIGCOMM*, 2009.
- [33] 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*, Nov 2007.
- [34] L. Gropok, D. Tse, and R. Yates, "Interference alignment for line-of-sight channels," in *arXiv:cs.IT/0809.3035*, 2008.
- [35] C. Suh and K. Ramchandran, "Exact regeneration codes for distributed storage repair using interference alignment," *CoRR*, vol. abs/1001.0107, 2010, <http://arxiv.org/abs/1001.0107>.
- [36] Y. Wu and A. Dimakis, "Reducing repair traffic for erasure coding-based storage via interference alignment," in *IEEE International Symposium on Information Theory*, July 2009, pp. 2276 –2280.
- [37] N. B. Shah, R. K. V., P. V. Kumar, and K. Ramchandran, "Explicit codes minimizing repair bandwidth for distributed storage," *CoRR*, vol. abs/0908.2984, 2009, <http://arxiv.org/abs/0908.2984>.
- [38] V. Cadambe, S. Jafar, and H. Maleki, "Optimal minimum repair bandwidth for exact regeneration in distributed storage," 2010.