

Exploiting Interference Locality in Coordinated Multi-Point Transmission Systems

Xinyu Zhang*, Mohammad A. Khojastepour[†], Karthikeyan Sundaresan[†],
Sampath Rangarajan[†], Kang G. Shin*

* The University of Michigan, [†] NEC Labs America

Email: xyzhang@eecs.umich.edu, {amir, karthiks, sampath}@nec-labs.com, kgshin@eecs.umich.edu

Abstract—Coordinated Multi-Point (CoMP) transmission is emerging as a concept that can substantially suppress interference, thus improving the capacity of multi-cell wireless networks. However, existing CoMP techniques either require sharing of data and channel state information (CSI) for all links in the network, or have limited capability of interference suppression. In this paper, we propose distributed interference alignment and cancellation (DIAC) to overcome these limitations. DIAC builds on a key intuition of interference locality — since each link is interfered with a limited number of neighboring links, it is sufficient to coordinate with those strong interferers and ignore others, in order to bound the overhead in CoMP. DIAC realizes the localized coordination by integrating interference cancellation and distributed interference alignment, and can be applied to both the uplink and downlink of multi-cell wireless networks. We validate DIAC using both model-driven and trace-based simulation where the traces are collected by implementing a MIMO-OFDM channel estimator on a software radio platform. Our experiments show that DIAC can substantially improve the degrees of freedom in multi-cell wireless networks.

I. INTRODUCTION

Recently, Coordinated Multi-Point (CoMP) transmission has emerged as a promising communication paradigm for multi-cell wireless networks. By allowing information sharing and cooperation among multiple transmitters, CoMP can dramatically reduce inter-cell interference, thereby improving the capacity of an entire network. Its advantages have been validated both theoretically [1] and experimentally [2].

CoMP requires the transmitters to share channel-state-information (CSI). By augmenting data sharing, interference between neighboring links can be further suppressed. For example, when using zero-forcing (ZF) precoding [3], multiple links are combined into one MIMO via CSI/data sharing between transmitters, and each link can transmit as if there were no mutual interference. However, such schemes need to group links into clusters. Links near the cluster edge still suffer from interference from neighboring clusters. To eliminate such edge effect and enable concurrent transmission of all links, the entire network of links need to be grouped and cooperate with each other. This becomes unrealistic for large networks, due to limited capacity of the backhaul network used for information sharing between links.

Interference alignment (IA) [4] is arguably the most promising mechanism that suppresses interference in CoMP systems. In theory, it can achieve $\frac{MK}{2}$ total degrees of freedom (DOFs) when there are K links in the network each with $M \times M$ MIMO, *i.e.*, half of the links can transmit concurrently. But this ideal bound has not been realized due to practical constraints,

and approximate solutions based on precoding and projection filtering have been proposed [5]. In [6], Gollakota *et al.* implemented a preliminary version of interference alignment called IAC which integrates IA with successive interference cancellation (SIC) on the uplink. IAC is centralized and is not scalable to large wireless networks, because it assumes all links interfere and share information with each other. Its achievable DOF is eventually limited by the number of antennas on each node, and does not scale as the network size grows.

Distributed algorithms with limited information sharing have also been proposed for CoMP. For example, NICE [7] allows uplinks that can decode their frames to send the decoded packets to other base stations (BSs), who then cancel such known interference using SIC. MaxSLNR [8] suppresses the leakage interference to neighboring links and only requires CSI sharing between neighbors. However, existing distributed algorithms for CoMP have limited capability for suppressing interference.

In this paper, we propose distributed interference alignment and cancellation (DIAC), which overcomes the above limitations by leveraging interference locality. Due to limited interference range of transmitters, each link is interfered by a limited number of neighboring links. So, it is sufficient to allow each link to coordinate with the neighboring links where dominant interference comes from. This in turn bounds the traffic load on the backhaul network. However, it is not straightforward to extend existing CoMP algorithms to support such *localized coordination*. For example, when using ZF precoding, the precoding vector designed for one link (and its neighbors) may conflict with that of a neighboring link (which may not have exactly the same set of neighbors). Consequently, the precoded signals may become random interference to neighboring links, and the resulting SINR may be even worse than without precoding. DIAC meets this challenge by a joint design of localized interference cancellation and distributed interference alignment.

In DIAC, each link shares data and CSI with, and cancels the interferences from only neighboring links (*i.e.*, a small set of coordination points). On the uplink, the base station receives decoded data from neighbors, reconstructs the interfering signals, and cancels them using successive interference cancellation (SIC). The downlink operation is more involved. Since clients cannot share data and perform SIC, the interference needs to be pre-cancelled at BSs. As mentioned above, existing pre-cancellation strategies for CoMP (*e.g.*, ZF-

precoding) become ineffective due to conflict of precoding vectors. DIAC solves this problem using a localized version of dirty-paper-coding (DPC), which only needs the composite DPC-coded information from neighbors. DIAC further applies distributed interference alignment (DIA) [5] that attempts to align interferences to a dimension orthogonal to the useful signals. This is realized by allowing links to share CSI with neighbors, and then iteratively design the precoding/projection matrix at the transmitter/receiver side, so as to minimize leakage interference to neighbors.

We evaluate the performance of DIAC using trace-driven simulation. We implement a MIMO-OFDM channel estimator on a WARP [9] software radio testbed, collect channel matrices between the transmitters and receivers, and then feed these traces to a Matlab simulator for DIAC. Our evaluation shows that DIAC can enable concurrent transmissions of multiple links even with a limited number of antennas. To further understand the performance of DIAC at large scale, we simulate DIAC under an empirical propagation model. Our experiments have revealed that the number of antennas and the scale of coordination may affect DIAC, but the effects diminish as both factors increase. Therefore, to exploit the benefits of DIAC while bounding the coordination overhead, it is sufficient for each link to restrict the coordination to a small set of neighboring links.

The remainder of this paper is structured as follows. In Sec. II, we present our system models and details of the DIAC algorithm. Sec. III, we introduce the design and implementation of the MIMO-OFDM channel estimator in the WARP platform. Sec. IV describes our experimental results and finally, Sec. V concludes the paper.

II. ALGORITHM DESIGN FOR DIAC

A. Basic models

1) *Precoding/projection model*: We consider a multi-cell network (wireless LANs or LTE cellular networks) consisting of K cells. Each cell includes one link (also referred to as one *user*) between a base station (BS) and a client. The AP and the client have M and N antennas, respectively. d streams of data are transmitted over each link, with $d \leq \min(M, N)$.

Let $X^{[k]}$ be the $d \times 1$ vector of data symbols to be transmitted by user k , V be the $M \times d$ precoding matrix for the transmitter, U the $N \times d$ projection matrix for the receiver, $H^{[kl]}$ the $M \times N$ channel matrix for transmitter l and receiver k . Using the superscript $[k]$ to denote variables for user k , the received signal of user k is:

$$Y^{[k]} = (U^{[k]})^\dagger \left(\sum_{l=1}^K H^{[kl]} V^{[l]} X^{[l]} \right) + Z^{[k]} \quad (1)$$

where $Z^{[k]}$ is the AWGN matrix; $(\cdot)^\dagger$ is the complex conjugate operator. We assume equal power allocation among data streams, *i.e.*, each stream has power of $\frac{P}{d}$, where P is the total power of a transmitter.

2) *Localized coordination*: CoMP systems may require coordinating base stations to share both data and CSI in real-time, and therefore, the backhaul wired network connecting them must have sufficiently large capacity and low latency to effectively support such coordination [10]. However, as the network size grows, such requirements become infeasible in practice.

In DIAC, we restrict the scale of coordination within the *locality* of each user, thereby bounding the amount of shared data/CSI and their propagation distance along the wired backhaul. We define the locality of coordination according to potential interference. The potential interference of link j to link i is:

$$I_{ij} = \frac{P_{ij}}{P_{ii}} \quad (2)$$

where P_{ij} is the signal power leaked from link j 's transmitter to link i 's receiver. A neighboring link j is included into link i 's set of coordinating points $R(i)$, only if I_{ij} is larger than a threshold T_i , *i.e.*,

$$R(i) = \{j : I_{ij} > T_i, j \neq i\} \quad (3)$$

As a result, each user i only needs to collect data/CSI from a small set of $|R(i)|$ neighbors, and suppress interferences from them.

Note that $j \notin R(i)$ may still cause interference to i , but DIAC does not attempt to suppress such interference, so that the locality of coordination ($|R(i)|$) can be bounded. $|R(i)|$ essentially reflects a tradeoff between the performance of DIAC and its overhead — a large $R(i)$ improves decoding performance by suppressing interference from more users, but meanwhile it requires the data and CSI to be shared with more users (*i.e.*, heavier load on the backhaul). We will evaluate such effects in our experiments (Sec. IV).

B. The DIAC Algorithm

In DIAC, we design a localized interference cancellation algorithm, and integrate it with distributed interference alignment (DIA), such that the network can support a large number of concurrent transmissions. In what follows, we first introduce the cancellation algorithm (for downlink and uplink, respectively), and then discuss how it can incorporate DIA.

1) *Challenges for localized cancellation*: Existing work has explored centralized algorithms for integrating interference cancellation with interference alignment in a fully-connected network (*i.e.*, every link interferes and coordinates with all other links) [6]. The basic idea is to first order the links, and ensure one link (*e.g.*, link 1) is decodable after interference alignment, and cancel its signals from link 2; and then cancel link 1 and link 2's signals at link 3, *etc.*. However, it is non-trivial to generalize it for practical partially-connected networks (*i.e.*, each link coordinates only with a limited number of neighboring links), because for such networks, no closed-form interference alignment solution exists that can ensure a specific link is decodable, and it is unknown how the links can be ordered to jointly perform cancellation and alignment.




Fig. 1. An example of localized downlink cancellation in DIAC.

The localized downlink interference cancellation is even more challenging and to our knowledge, has not been explored in previous works. Intuitively, the downlink cancellation can be realized by joint precoding between neighboring BSs. However, in partially-connected networks, joint precoding becomes infeasible due to conflicting precoding vectors between different neighboring sets. For example, this happens in Fig. 1 when locality of coordination is restricted to adjacent links. It is feasible for BS2 to jointly design the precoder with BS1 and BS3, and BS3 with BS2 and BS4. But the resulting precoding vectors of BS2 and BS3 will conflict with each other, as they take into account information from different sets of links. Thus a straightforward application of the precoding may cause random interference, and perform even worse than without precoding.

2) *Localized downlink interference cancellation*: On the downlink, each user i *pre-cancels* interference from neighbors in $R(i)$ in a distributed manner. Simply put, whenever a user $j \in R(i)$ is decodable (*i.e.*, it has sufficient SINR to decode its data), its BS shares the to-be-sent data (interference for other users) with users $\{k : j \in R(k)\}$ (which includes i). Then BS i encodes and pre-cancels the interference using dirty-paper-coding (DPC) [3], based on the channel matrix from the transmitters in $R(i)$ and its own receiver. Afterwards, the receiver of link i no longer experiences any interference from the transmitter of link j . Note that DIAC also requires the channel matrix between dominant interferers (those in $R(i)$) and the receiver of i . The estimation of channel matrix between each transmitter and receiver will be discussed in Sec. III.

A key feature of the downlink cancellation in DIAC is decentralization: it only requires neighboring users to exchange encoded signals, without accounting for other interferers. Fig. 1 illustrates an example of such localized downlink cancellation, where we restrict the locality of coordination to adjacent users. Suppose user 1 is able to decode its packet (possibly because of shorter link distance and higher SINR), then the corresponding transmitter (BS1) will pass the encoded samples $X^{[1]}$ of each frame to the transmitter of user 2 (*i.e.*, BS2) through the wired backhaul. BS2 then encodes these samples using DPC in order to pre-cancel them. The resulting samples sent by BS2 will be a function of $X^{[1]}$ and its own data samples $X^{[2]}$. We denote this operation as: $\chi^{[2]} \leftarrow \text{DPC}(X^{[1]}, X^{[2]})$. After cancelling interference from user 1, user 2 may have sufficient SINR to decode its own frames (despite the unknown interference from user 3). Then, BS2 will pass its signals $\chi^{[2]}$ to the neighboring user 3 who cannot decode its signals yet. This procedure continues until no further pre-cancellation is possible. Afterwards, the users

Algorithm 1 Downlink operations of each user i in DIAC.

1. **Input**: vector of data $X^{[i]}$; receiver projection matrix $U^{[i]}$ and transmitter precoding matrix $V^{[i]}$.
/ $U^{[i]}$ is calculated from Eq. (6) to Eq. (7); $V^{[i]}$ is calculated as its dual (Sec. II-B4) */*
 2. */* BS operations: */*
 3. $\chi^{[i]} = V^{[i]}X^{[i]}$ */* Precoding for interference alignment */*
 4. **while** link i is not decodable and received no sync
 5. **do**
 6. Receive shared data $\chi^{[k]}$ from the BS of $\{k : k \in R(i), k \text{ is decodable}\}$.
 7. Encode to precancel the interference:
 $\chi^{[i]} \leftarrow \text{DPC}(\chi^{[i]}, \chi^{[k]})$
 8. **if** No update within the current round
 9. BS i sends notification message to coordinator.
 10. **endif**
 11. **done**
 12. **if** link i is decodable
 13. pass $\chi^{[i]}$ to the BS of $\{j : j \in R(i), j \text{ is not decodable}\}$
 14. BS i waits for sync message and then send $\chi^{[i]}$.
 15. **endif**
 16. */* Client's operations: */*
 17. After reception, client i applies interference projection:
 $Y^{[i]} \leftarrow U^{[i]}Y^{[i]}$ and then decode.
 18. */* Coordinator's operations: */*
 19. **foreach** time slot
 20. **if** received notification message from all BSs
 21. Broadcast sync message
 22. **endif**
 23. **endifor**
-

can transmit concurrently. Although the downlink interference cancellation is restricted within $R(i)$, we will show in Sec. IV it can still substantially improve the number of concurrent transmissions.

3) *Localized uplink interference cancellation*: The uplink differs from the downlink in the content of information sharing — only decoded data needs to be shared between neighboring links. Take the network in Fig. 1 as an example. When BS1 can decode the data from client1, it will pass the digital data $X^{[1]}$ sent by client1 to BS2. BS2 then cancels the interferences from client1 using successive interference cancellation (SIC). More specifically, it needs to reconstruct the interfering signals based on $X^{[1]}$ and the channel matrix from client1 to BS2, and then subtract them from the composite signals ($Y^{[2]}$) it received. Afterwards, if BS2 is able to decode its data, it will pass the original data $X^{[2]}$ in the same way (instead of passing a function of the data from both BS1 and BS2, as in the downlink case).

4) *Matrix design for distributed interference alignment*: Besides the localized cancellation, DIAC employs distributed interference alignment (DIA) [5], [11], which calculates the precoding/projection matrix, so as to align interferences into

a space orthogonal to useful signals. For the sake of completeness, we briefly introduce the matrix design for DIA, and then describe how it can be integrated with the localized cancellation algorithms in Sec. II-B5.

In classical interference alignment, all interferers' signals are aligned into the null space of each receiver k 's projection matrix [5]. Equivalently,

$$(U^{[k]})^\dagger H^{[kj]} V^{[j]} = 0, \forall j \neq k \quad (4)$$

This is often an over-constrained system of equations, and the closed-form solution remains an open problem. However, DIA [5] can be used as an approximate solution to the interference alignment problem, and is guaranteed to converge. The basic idea is to iteratively design the matrix U and V to minimize the LHS of Eq. (4). In each iteration, each receiver i calculates the projection matrix $U^{[i]}$ to minimize the leakage to $j : i \in R(j)$:

$$U^{[i]} = \arg \min_{U^{[i]}} I^{[i]} = \text{Tr} \left[(U^{[k]})^\dagger Q^{[i]} U^{[k]} \right] \quad (5)$$

$$\text{where } Q^{[i]} = \sum_{j \in R(i)} \frac{P}{d} H^{[ij]} V^{[j]} (V^{[j]})^\dagger (H^{[ij]})^\dagger \quad (6)$$

is the interference covariance matrix at receiver k . The solution to the minimization problem is [5]:

$$U^{[k]} = u_d[Q^{[k]}] \quad (7)$$

where $u_d[A]$ denotes the d eigenvectors corresponding to the d smallest eigenvalues of matrix A . Then, given U , DIAC designs the precoding matrix V in a similarly way, but reverses the roles of the transmitter and receiver. The procedure iterates until the leakage interference $I^{[i]}$ is below a small threshold. The resulting precoding/projection matrices are used as input to the localized cancellation algorithms.

5) *Summary of DIAC*: Based on the above description of operations, we summarize the downlink DIAC in Algorithm 1.

On the downlink, before applying DPC, the BS must precode the information bits (line 3) and then use the precoded data as input to the cancellation procedure. We assume the data/CSI sharing (before the actual transmission) can be done in a time-slotted, round-based manner. In each slot, each BS performs decodability check, information sharing, and precancellation. If no operation can be done in the current time slot, it sends a notification message to a central coordinator (which can be one of the base stations). Once the coordinator receives the notification message from all users within the same slot, this time slot is used as a synchronization barrier — the coordinator will send a sync message to all BSs, allowing them to transmit concurrently. Operations (*e.g.*, channel estimation and CSI feedback) before the synchronized transmission can be realized by traditional medium access mechanisms, such as TDMA or OFDMA.

The uplink operations can be done in a similar manner (we omit the details to avoid duplication). Note that DIAC requires the clients to send frames concurrently. We assume this is realized by synchronizing all BSs, and then allowing each BS to send a sync message to its client. The cancellation procedure is performed in a time-slotted manner similar to the uplink,

and the coordinator will broadcast a *stop* message if it receives a notification message from every BS.

6) *SINR analysis*: After obtaining the V and U matrix, the SINR of the l -th stream of k -th receiver is:

$$\text{SINR}_{kl} = \frac{(U_{*l}^{[k]})^\dagger H^{[kk]} V_{*l}^{[k]} (V_{*l}^{[k]})^\dagger (H^{[kk]})^\dagger U_{*l}^{[k]} P}{(U_{*l}^{[k]})^\dagger B^{[kl]} U_{*l}^{[k]} P} \frac{P}{d} \quad (8)$$

where

$$B^{[kl]} = \sum_{j \in R(k)} \frac{P}{d} \sum_{e \neq l} H^{[kj]} V_{*e}^{[j]} (V_{*e}^{[j]})^\dagger (H^{[kj]})^\dagger + I_{N^{[k]}} \quad (9)$$

is the interference covariance matrix. $U_{*l}^{[k]}$ denotes the l -th column of matrix $U^{[k]}$.

Assuming interferences are perfectly cancelled, after cancelling interference from user i , the resulting interference covariance matrix becomes:

$$B^{[kl]} = \sum_{j \in R(k), j \neq i} \frac{P}{d} \sum_{e \neq l} H^{[kj]} V_{*e}^{[j]} (V_{*e}^{[j]})^\dagger (H^{[kj]})^\dagger + I_{N^{[k]}} \quad (10)$$

III. IMPLEMENTATION OF MIMO-OFDM CHANNEL ESTIMATOR

Our implementation consists of two parts: Matlab based simulation of the DIAC algorithm, and software radio (WARPLab [9]) based implementation of MIMO channel estimation algorithm, which is used to obtain traces of channel matrix between multiple links.

Fig. 2 and Fig. 3 show the transmitted packet format for channel estimation, and the corresponding detection and estimation procedure at the receiver. The MIMO-OFDM packet preamble format follows that of the 802.11n standard [12]. An 802.11 preamble (also referred to as *STF*) has 160 complex samples and occupies 64 subcarriers. From time-domain perspective, the STF is a periodic signal that repeats every 16 complex samples [13]. The receiver performs *self-correlation* between the latest 16 samples and previous 16 samples, which has an outstanding output only if two consecutive sequences of samples match (*i.e.*, an STF appears), and the corresponding output is comparable to the signal's energy level [13]. After detecting the STF, the receiver further performs *cross-correlation* between the received STF samples and the original samples in the STF. An outstanding peak appears only when the received samples align with the known STF, and the peak position is used as a synchronization point marking the start of the packet.

To estimate the channel coefficients (including magnitude attenuation and phase distortion) of each subcarrier and the frequency offset between transmitter and receiver, an additional preamble is used, called *long-training field* (LTF). LTF comprises two duplicated versions of a random sequence (consisting of 1 and -1) of length 64 carried by the 64 subcarriers [14]. To obtain the channel coefficients and frequency offset, the receiver performs self-correlation between the two truncated random sequences and normalizes it by the magnitude, similar to an 802.11 channel estimator [13].




Fig. 2. Packet format for channel estimation.




Fig. 3. Packet detection and channel estimation procedure.

When running MIMO, all the transmit antennas share the same STF, but their LTFs are transmitted sequentially, *i.e.*, while one antenna is transmitting, all others are silent (transmitting zero-power signals). Meanwhile, each receive antenna can estimate the channel between the active transmit antenna and itself.

Our ongoing work involves a full-fledged implementation of DIAC on WARP. As DIAC requires synchronization between neighboring base stations, we have modified the FPGA module in WARP to synchronize the carrier frequency, sampling clock, and packet starting time between multiple MIMO transmitters. We plan to implement a real-time version of DIAC on this platform, and demonstrate its performance in real environment.

IV. EXPERIMENTAL RESULTS

We evaluate DIAC using *i)* trace-driven simulation, which collects real channel traces from a WARP MIMO testbed. *ii)* model-driven simulation, which applies an empirical pathloss and fading model to a synthetic topology. As benchmark comparison, we also consider the following distributed algorithms that can be applied to partially-connected networks, where each link either transmit independently, or only needs to collect data/CSI from neighboring links that interfere with it. *i)* Single-user MIMO beamforming via SVD [3]. *ii)* Maximizing signal to leakage-interference and noise ratio (maxSLNR) [8]. *iii)* Distributed interference alignment (DIA) [5], [11], which is equivalent to the alignment iteration of DIAC, *i.e.*, each user attempts to align interference from others without cancellation.

A. Trace based simulation

We set up a WARP testbed consisting of four 4×2 MIMO links located in an office environment. Fig. 4 shows the node locations. Two links are located within two different offices separated by two walls and a corridor; the other two links are located in the corridor. Due to the short transmission range of WARP radios, the distance between each transmitter and receiver is 2 to 5 meters; and the distance between each transmitter and neighboring transmitters is 4 to 8 meters. Given this topology, we test the DIAC and benchmark algorithms based on the MIMO-OFDM channel estimation results.

Fig. 6(a) shows the resulting SINR distribution when each link uses 4×2 MIMO to send one data stream. The interference




Fig. 4. Testbed node locations.

Fig. 5. Model based simulation results.

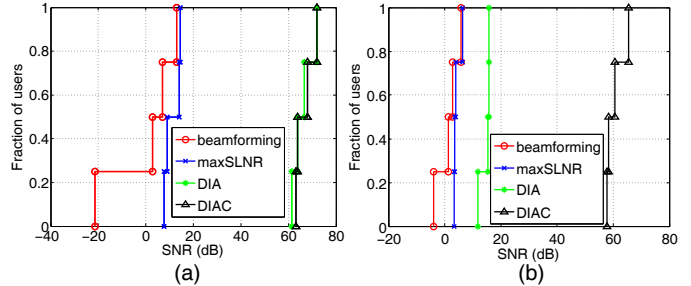


Fig. 6. Trace based simulation results with (a) 4×2 MIMO and (b) 2×2 MIMO

alignment based approaches (*i.e.*, DIA and DIAC) outperform beamforming and maxSLNR by around 50dB, implying that they may completely suppress interference and enable concurrent transmissions, whereas alternative approaches result in collisions. Notably, DIAC shows marginal gain over DIA under this setting. This is because the 4×2 MIMO is sufficient for DIA to suppress all interferences even without cancellation. When using 2×2 antennas, however, DIA alone is unable to achieve perfect alignment. All 4 links have less than 18dB SINR (Fig. 6(b)). In contrast, by integrating interference cancellation, DIAC achieves more than 40dB gain over DIA, thereby maintaining concurrent transmission of all links.

B. Model-driven simulation

Due to a limited number of WARP nodes, the trace-driven simulation only involves 4 MIMO links. To fully understand the factors that affect DIAC's performance in general network topologies, we further conduct model-driven simulation. We use an empirical propagation model recommended by the IEEE 802.15 for 2.4GHz indoor environment [15] to model large-scale fading. At distance d , the signal's path-loss (in dB) is:

$$L_{dB}(d) = \begin{cases} 40.2 + 20 \log_{10}(d), & d \leq 8m \\ 58.5 + 33 \log_{10}(d/8), & d > 8m \end{cases}$$

For small-scale fading, we use the Rayleigh fading model. When running DIAC, we assume the SNR threshold for decoding is 10 dB. The transmit power of each node is $P_t = 15\text{dBm}$. We further assume the receiver noise power is 10^{-8} of the transmit power level.

We first generate a line topology in which 16 APs are located on a straight line, with 30m separation. Each AP has one client, randomly located within a circle (radius 15m) around it. For simplicity, we define locality of coordination using hop-distance for this topology. We set the locality to 4, *i.e.*, each link coordinates (share data and CSI) with




Fig. 7. Effect of locality.




Fig. 8. Effect of number of antennas.

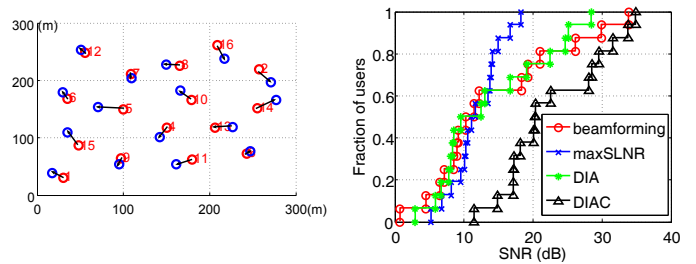


Fig. 9. Node locations in a random topology.

Fig. 10. SNR distribution in a random topology.

neighbors within 4 hops. Fig. 5 shows the resulting SNR distribution of all 16 links with 2x2 antennas and single-stream transmission. In this network, each link may be interfered by many other links. Therefore, when using DIA, the 2x2 MIMO is insufficient to align all interferences to a space orthogonal to the desired signals of each link. Consequently, DIA only performs slightly better than beamforming. maxSLNR outperforms beamforming for low-SINR links that suffer from severe interference, but has lower performance for high-SINR links because it dissipates part of the power to suppressing leakage interference, instead of focusing all power towards the intended receiver. Again, by combining cancellation with alignment, DIAC outperforms all other algorithms by around 10dB for all links.

Fig. 7 shows the effect of locality on the mean SINR of all links when using DIAC. When locality is small, SINR increases quickly with locality. However, further improving locality to large numbers gains marginal improvement, mainly because links further away cause negligible amount of interference. Therefore, even if we localize the coordination scale to a small number of neighbors, DIAC is still able to enable concurrent transmission of all links.

Fig. 8 shows the effect of the number of antennas on mean SINR. maxSLNR outperforms beamforming only when a large number of antennas are used to suppress interference. Both schemes have around 20dB lower SINR than DIA and DIAC. When using a small number of antennas, DIA may have much lower SINR than DIAC, implying that interference alignment highly depends on the available antennas. By combining cancellation, DIAC is able to maintain high SINR for all links even with a small number of antennas.

We further evaluate DIAC in a random topology resembling real-world multi-cell wireless LANs. As shown in Fig. 4, the base stations are randomly generated within a 300m x 300m field, but neighboring base stations are separated by at least 30m. Then a client location is randomly generated within a radius of 10m around each BS. The resulting SINR distribution is plotted in Fig. 10, where we configure locality according to $T_i = -45dB, \forall i$. We observe that DIA may perform even worse than beamforming, mainly because in this topology, each link suffers from a larger number of interferers than in the line topology. Whereas with DIAC, SINR is above 15dB for most of the links, implying that DIAC is able to achieve a high level of concurrent transmissions even in real-world network topologies.

V. CONCLUSION

In this paper, we have proposed DIAC, a distributed algorithm that improves the performance of CoMP systems by leveraging the locality of interference. DIAC integrates localized interference cancellation algorithms with distributed interference alignment, is applicable to both uplink and downlink transmissions, and only requires data/CSI sharing between close-by neighboring base stations. Both trace-driven simulation and model-driven simulation have shown DIAC to substantially improve the number of concurrent transmissions, even with a limited number of antennas.

REFERENCES

- [1] Sawahashi, M. and Kishiyama, Y. and Morimoto, A. and Nishikawa, D. and Tanno, M., "Coordinated Multipoint Transmission/Reception Techniques for LTE-Advanced [Coordinated and Distributed MIMO]," *IEEE Wireless Communications*, vol. 17, no. 3, 2010.
- [2] V. Jungnickel, A. Forck, S. Jaeckel, F. Bauermeister, S. Schiffermueller, S. Schubert, S. Wahls, L. Thiele, T. Haustein, W. Kreher, J. Mueller, H. Droste, and G. Kadel, "Field Trials Using Coordinated Multi-Point Transmission in the Downlink," in *IEEE PIMRC Workshops*, 2010.
- [3] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
- [4] V. Cadambe and S. Jafar, "Interference Alignment and Degrees of Freedom of the K-User Interference Channel," *IEEE Transactions on Information Theory*, vol. 54, no. 8, 2008.
- [5] K. S. Gomadam, V. R. Cadambe, and S. A. Jafar, "Approaching the Capacity of Wireless Networks through Distributed Interference Alignment," *CoRR*, vol. abs/0803.3816, 2008.
- [6] S. Gollakota, S. D. Perli, and D. Katabi, "Interference Alignment and Cancellation," in *Proc. of ACM SIGCOMM*, 2009.
- [7] K. Balachandran, J. Kang, K. Karakayali, and K. Rege, "NICE: A Network Interference Cancellation Engine for Opportunistic Uplink Cooperation in Wireless Networks," *IEEE Transactions on Wireless Communications*, vol. 10, no. 2, 2011.
- [8] M. Sadek, A. Tarighat, and A. Sayed, "A Leakage-Based Precoding Scheme for Downlink Multi-User MIMO Channels," *Wireless Communications, IEEE Transactions on*, vol. 6, no. 5, 2007.
- [9] A. Khattab, J. Camp, C. Hunter, P. Murphy, A. Sabharwal, and E. W. Knightly, "WARP: a Flexible Platform for Clean-Slate Wireless Medium Access Protocol Design," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 12, 2008.
- [10] D. Gesbert, S. Hanly, H. Huang, S. Shamai Shitz, O. Simeone, and W. Yu, "Multi-Cell MIMO Cooperative Networks: A New Look at Interference," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 28, no. 9, 2010.
- [11] M. Guillaud and D. Gesbert, "Interference Alignment In The Partially Connected K-User MIMO Interference Channel," in *Proc. of European Signal Processing Conference (EUSIPCO)*, 2011.
- [12] IEEE Standard, "802.11n: Enhancements for Higher Throughput," 2009.
- [13] K. Jamieson, "The SoftPHY Abstraction: from Packets to Symbols in Wireless Network Design," Ph.D. Thesis, MIT, 2008.
- [14] IEEE Standard, "802.11TM: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," 2007.
- [15] IEEE 802.15 Working Group, "Coexistence Analysis of IEEE Std 802.15.4 With Other IEEE Standards and Proposed Standards," 2010.