

MIT Open Access Articles

Interference alignment by motion

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Adib, Fadel, Swarun Kumar, Omid Aryan, Shyamnath Gollakota, and Dina Katabi. "Interference Alignment by Motion." Proceedings of the 19th Annual International Conference on Mobile Computing & Networking - MobiCom '13 (2013), September 30–October 4, Miami, FL, USA.

As Published: <http://dx.doi.org/10.1145/2500423.2500449>

Publisher: Association for Computing Machinery

Persistent URL: <http://hdl.handle.net/1721.1/87093>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Interference Alignment by Motion

Fadel Adib^{†1} Swaran Kumar^{†1} Omid Aryan¹ Shyamnath Gollakota² Dina Katabi¹
¹Massachusetts Institute of Technology {fadel, swaran, omida, dk}@mit.edu ²University of Washington gshyam@cs.washington.edu
[†]Co-primary authors

ABSTRACT

Recent years have witnessed increasing interest in interference alignment which has been demonstrated to deliver gains for wireless networks both analytically and empirically. Typically, interference alignment is achieved by having a MIMO sender precode its transmission to align it at the receiver. In this paper, we show, for the first time, that interference alignment can be achieved via motion, and works even for single-antenna transmitters. Specifically, this alignment can be achieved purely by sliding the receiver's antenna. Interestingly, the amount of antenna displacement is of the order of one inch which makes it practical to incorporate into recent sliding antennas available on the market. We implemented our design on USRPs and demonstrated that it can deliver $1.98\times$ throughput gains over 802.11n in networks with both single-antenna and multi-antenna nodes.

Categories and Subject Descriptors C.2.2 [Computer Systems Organization]: Computer-Communications Networks

Keywords Interference Alignment, Interference Nulling, Sliding Antennas, Motion-based Interference Management, MIMO, Wireless

1. INTRODUCTION

Recent advances in MIMO technology have demonstrated the practicality of interference alignment and its throughput gains [7, 26, 15, 13]. These systems leverage the MIMO capability of the transmitter to precode the signal and align it along a particular spatial direction at the receiver. In contrast, this paper investigates a simple question: Can we perform interference alignment with single-antenna transmitters? Furthermore, can interference alignment be performed purely by the receiver without cooperation from the senders? A positive answer to these questions could extend the gains of interference alignment to scenarios with single antenna systems, such as sensors, hence improving the overall throughput of these networks. Furthermore, the ability to do alignment purely at the receiver eliminates the need to send feedback from the receiver to the senders to inform them about the direction along which they should align their signals.

Motivated by the above questions, we investigate whether a receiver can perform interference alignment by simply adjusting the position of one of its antennas. Our intuition is that this may be pos-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom'13, September 30–October 4, Miami, FL, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-1999-7/13/09

<http://dx.doi.org/10.1145/2500423.2500449> ...\$15.00.

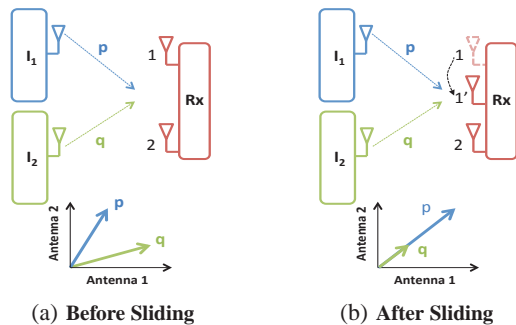


Figure 1: Interference Alignment by motion. Consider two single-antenna interferers I_1 and I_2 and a two-antenna receiver R_x . Initially, the two interferers are not aligned in the antenna space of R_x . Then, R_x slides one of its receive antennas so that I_1 and I_2 are aligned in the antenna space.

sible due to two reasons. First, indoor settings are rich in multipath. Hence, at any point in space, the perceived channel is a combination of many channels that correspond to the various paths that the signal traverses. Even a small shift in the antenna location would cause some of these paths to become shorter and others longer, leading to a significant change in the resulting channel. Second, channels are continuous functions over space.¹ This means that the receiver can gradually slide its antenna, measure the channels, and keep moving the antenna in the direction that increases the alignment between the senders.

To test this intuition, we used USRP radios to experiment with the scenario in Fig. 1. In this setup, we have two single-antenna interferers and a two-antenna receiver. We mount one of the receive antennas on an iRobot Create robot to emulate a sliding antenna. We measure the channels and gradually adjust the position of one of the receive antennas, moving it in the spatial direction that increases the alignment between the interferers, as in Fig. 1(b). The experiments have confirmed our intuition: alignment can indeed be performed by sliding the receiver's antenna. What is perhaps even more surprising is that the required antenna displacement was typically an inch or less. Multiple experiments in two campus buildings in both line-of-sight (LOS) and non-line-of-sight (NLOS) scenarios showed that in 84% of the settings, the required antenna displacement was less than one inch. The implications of our results are two-fold: First, they show for the first time that interference alignment can be achieved via motion. Second, the fact that the required displacement is of the order of one inch, means that motion-based interference alignment can be incorporated into recent sliding antennas available on the market.²

To gain insight into these empirical results, we have extended past signal propagation models to incorporate the effects of multi-

¹The channel is a linear combination of multiple continuous waveforms [27], and, hence, is continuous over space and time.

²Recent USB Wi-Fi adapters have sliding antennas whose positions can be manually adjusted to improve the SNR [28].

path on interference alignment. Our results reveal that the presence of a reflector near the receiver creates significant spatial diversity, causing the channel to change dramatically (from combining constructively to combining destructively) within an inch or two. Since the receiver is typically situated on some platform which serves as a reflecting surface (e.g., a table, the floor, etc.), it is likely that it experiences reflections from at least one nearby reflector, enabling motion based interference alignment.

We introduce MoMIMO, a technique that enables MIMO interference management through mobility. Besides motion-based interference alignment, our study of MoMIMO reveals several other interesting capabilities:

- First, MoMIMO also provides interference nulling. Specifically, a single-antenna receiver can slide its antenna position to cancel an unwanted interference signal from a single-antenna transmitter, hence enabling single-antenna nodes to exploit MIMO nulling.
- Second, the process of sliding antennas to seek positions of interference alignment or nulling can be automated using a Stochastic Hill Climbing algorithm, as described in §6.
- Third, both uplink and downlink traffic benefit from motion-based MIMO techniques. This stems from the reciprocity of wireless channels, which ensures that motion-based alignment or nulling on the downlink, causes interference alignment or nulling on the uplink (see §7).

We implemented MoMIMO on an indoor network of USRP N210 software radios. To emulate a sliding antenna, we mounted the antennas on an iRobot Create robot. We conducted our experiments at a bandwidth of 20 MHz in the 2.4 GHz Wi-Fi band, both in line-of-sight and non-line-of-sight scenarios. Our results show the following:

- MoMIMO reduces the interference-to-noise ratio (INR) of undesired signal by an average of 22 dB for interference alignment and 15 dB for interference nulling.
- In a network of multiple transmitter-receiver pairs that use a combination of alignment and nulling, MoMIMO achieves an average throughput gain of $1.98\times$ over 802.11n in networks with both single-antenna and multi-antenna nodes.
- MoMIMO’s Stochastic Hill Climbing algorithm requires a mean displacement of the receive antenna of 0.44 inches for interference alignment and 1.17 inches for nulling.³
- Due to reciprocity, MoMIMO’s interference alignment and nulling on the downlink also reduces the INR on the uplink to below the noise floor (i.e., the mean INR on the uplink is below 0 dB).
- MoMIMO is robust to sudden channel changes in otherwise static environments. It adopts a fail safe mechanism that falls back to 802.11n if the channel changes significantly, and recovers its gains once the channel is stable.

Contributions: This paper has three main contributions: (1) It is the first to demonstrate the feasibility of interference alignment and nulling purely via motion. Furthermore, it reveals that the required displacement is about one inch, and hence can potentially be supported by sliding antennas similar to those used in recent products [28]. (2) It presents a stochastic hill climbing algorithm that automatically repositions the antenna to achieve interference alignment or nulling. (3) The paper implements its design and demonstrates large throughput gains over real wireless channels.

³Alignment is easier to satisfy because many options exist for the direction along which two senders may be aligned, which provides extra flexibility for the choice of channels.




Figure 2: Interference Nulling. The transmitter cancels its signal at Rx_2 by precoding it, while Rx_1 can still receive its signal x .




Figure 3: Interference Alignment. q 's transmitter performs precoding in order to align q along the same direction as p in the antenna space of the receiver.

2. BACKGROUND

We briefly introduce interference nulling and alignment.

(a) Interference Nulling: Interference nulling enables a MIMO node to transmit signals, and at the same time cancel them out at particular receiver (other than its own). Consider the example in Fig. 2; it consists of a two-antenna transmitter Tx and two single-antenna receivers Rx_1 and Rx_2 . Let h_1 and h_2 denote the wireless channels from the two transmit antennas of Tx to the single antenna of Rx_2 . Suppose the transmitter wants to send a signal x to Rx_1 without interfering at Rx_2 . The transmitter can then send αx on its first antenna and βx on its second antenna. Then, Rx_2 would receive: $y = (\alpha h_1 + \beta h_2)x$. Hence, the transmitter can null its signal to this receiver, simply by picking $\alpha = -h_2$ and $\beta = h_1$. This process of encoding the transmitted signal across multiple antennas (i.e. multiplying by α and β) is referred to as *precoding*, and is used in both nulling and alignment.

(b) Interference Alignment: Recall the following basic properties of MIMO systems [27, 15]:

- An N -antenna receiver receives signals on each of its N antennas. These signals can be represented as one vector in an N dimensional space.
- An N -antenna receiver can decode up to N independent signals.

Suppose a 2-antenna receiver receives two signals – a desired signal x and an interfering signal p . As mentioned above, these signals can be represented as vectors in a 2-D space, and the MIMO receiver can decode them. However, let’s say another interferer joins the network, so that the 2-D space has a third vector q , as in Fig. 3(a). Unfortunately, the receiver can now no longer decode. So, how can the transmitter of q avoid interfering at the receiver?

The transmitter of q can precode its transmission to rotate the vector q and align it along the same direction as p ; this technique is called *interference alignment*.⁴ In effect, the receiver now obtains only two vectors, one vector is the desired signal x , and the other is the sum of interferences p and q , as shown in Fig. 3(b). Now, the receiver can eliminate all interference by projecting the received

⁴Specifically, to align the vector q along the vector p , the transmitter uses a precoding matrix M , such that, $p^\perp qM = 0$, where p^\perp denotes the vector space orthogonal to p [15, 7].

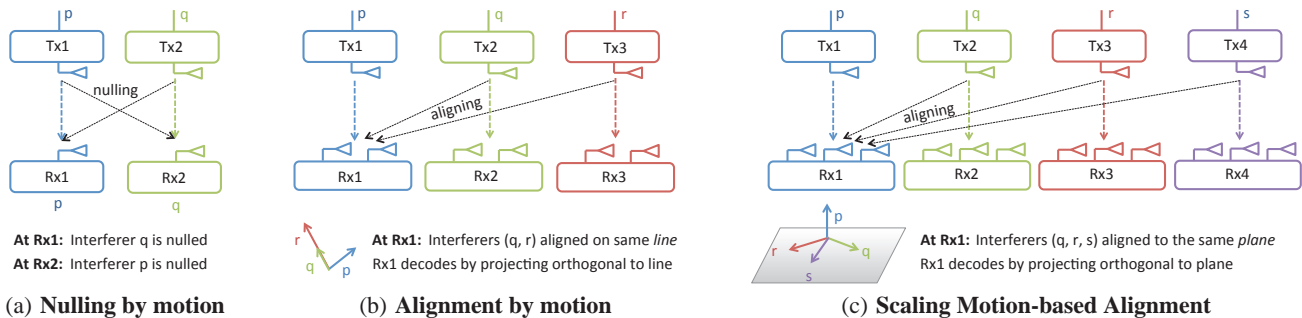


Figure 4: MoMIMO enables concurrent transmissions. The figure shows three examples: **(a) Nulling:** Rx1 slides its antenna to null interference from Tx2, and Rx2 slides its antenna to null interference from Tx1. Hence, the network supports 2 concurrent transmissions. **(b) Alignment:** Rx1 slides its antenna to align interference from Tx2 and Tx3. Similarly, Rx2 and Rx3 slide their antennas to align their interferers. Thus, the network can support 3 concurrent streams. **(c) Scaling Alignment:** Rx1 slides one of its antennas to align interference from Tx2, Tx3 and Tx4. Similarly, Rx2, Rx3, and Rx4 each slide one of their antennas to align their interferers. Thus, the network supports 4 concurrent streams.

signal on a direction orthogonal to the direction of the aligned interference.

3. SCOPE OF MoMIMO

MoMIMO enables wireless nodes to perform techniques such as interference alignment and nulling by sliding one of their antennas by an inch or two. MoMIMO applies to both single antenna and multi-antenna nodes, i.e., MoMIMO enables both single and multi-antenna nodes to manage interference in a manner that increases the number of concurrent streams beyond what is available purely by the number of antennas on these nodes.

MoMIMO is not meant for mobile or hand-held devices, but is instead targeted for relatively static environments, with stable interference patterns. Examples of such scenarios include: (1) A home Wi-Fi access point that suffers interference from the neighboring apartment. (2) A sensor network collection point suffering from interference. (3) Wireless as a replacement for wires at homes and offices, e.g., a Wi-Fi link used to connect a TV to a DVD player [24], wireless surveillance cameras [3], etc. (4) Flyways in Data centers [11].

4. MoMIMO IN A NETWORK

We begin by presenting examples that illustrate how MoMIMO can provide significant throughput gains for wireless networks. We explain when nodes may use motion-based alignment and when they use motion-based nulling. For this section, we assume that it is possible to find locations that satisfy the desired interference alignment or interference nulling, by using small position adjustments of an inch or two. We will validate this assumption both empirically in §9 and analytically in §5.

4.1 Motion-Based Interference Nulling

Consider two co-located single-antenna transmitter-receiver pairs: Tx1-Rx1 and Tx2-Rx2 (Fig. 4(a)). Say these nodes are initially at positions where they interfere strongly (i.e. the interference to noise ratio (INR) from Tx1 to Rx2 and Tx2 to Rx1 is high). In such a scenario, 802.11 permits only one of Tx1 and Tx2 to transmit at any given time, i.e., the maximum number of concurrent streams is one.

Recall that because Tx1 and Tx2 are single-antenna transmitters, they cannot perform standard interference nulling to each other's receivers. This is because standard interference nulling requires a node to precode its signal across *multiple transmit antennas* as discussed in §2.

In contrast, MoMIMO can leverage motion-based nulling to send two concurrent streams that do not interfere with each other. Specif-

ically, each receiver makes small adjustments to its antenna position by one or two inches until the signal from the interfering transmitter is nulled (i.e. below the noise floor). In this example, Rx1 adjusts its antenna until it nulls Tx2's interference. Similarly, Rx2 adjusts its antenna until it nulls Tx1. Now that there is no interference between the two streams, Tx1 and Tx2 can transmit concurrently to their respective receivers.

Reciprocity: We now have the receive antennas at positions where Tx1 and Tx2 can transmit concurrently, without interference. But what happens when Rx1 and Rx2 need to transmit their acknowledgments concurrently to Tx1 and Tx2? Interestingly, we notice that there is no need for any new adjustments in antenna positions, i.e. the ACKs do not interfere at Tx1 and Tx2. This is because reciprocity in wireless channels also leads to reciprocity in motion-based nulling (details in §7). Thus, once receivers perform motion-based nulling, they not only receive their desired signals concurrently, but also deliver ACKs to their transmitters, interference-free.

Multiplexing Gains: MoMIMO's motion-based interference nulling, can enable two concurrent streams in this setting, i.e., a $2\times$ multiplexing gain over 802.11.

4.2 Motion-Based Interference Alignment

Consider three co-located transmitter-receiver pairs: Tx1-Rx1, Tx2-Rx2 and Tx3-Rx3, as shown in Fig. 4(b)). The transmitters have only one antenna each, while the receivers are two-antenna nodes. Once again, these nodes are in positions where they interfere strongly with each other. Therefore, 802.11 permits only one of Tx1, Tx2, and Tx3 to transmit packets at a given point in time; hence, the maximum number of concurrent streams is one.

Recall that since Tx1, Tx2, and Tx3 have only one antenna each, they cannot align their interfering signals along a common spatial direction at each receiver. This is because such alignment requires the transmitter to precode its signal across *multiple transmit antennas* as explained in §2.

In contrast, MoMIMO leverages motion-based alignment to send three concurrent streams that do not interfere with each other. Let's consider the two-antenna receiver Rx1, which receives its desired signal from Tx1 and unwanted interference from both Tx2 and Tx3. To mitigate this interference, Rx1 performs motion-based alignment by adjusting one of its antennas so that the signals from its two interfering transmitters, Tx2 and Tx3, are aligned in its 2-D space (see Fig. 4(b)). Now Rx1 obtains only two signal vectors in this 2-D space: the signal from its desired transmitter (Tx1) and the sum of the interference from the unwanted transmitters Tx2 and Tx3. Thus, as in standard interference alignment (see §2), the receiver can eliminate all interference by projecting its received sig-




Figure 5: MoMIMO in heterogeneous networks. The two single antenna pairs perform nulling as in Fig. 4(a). Rx3 slides one of its antennas to align Tx1 and Tx2; similarly, Tx3 slides one of its antennas to align ACKs from Rx1 and Rx2. Tx3 also precodes its signal to null at Rx1. The network can support 3 concurrent streams.

nal along a direction orthogonal to that of the aligned interference. Similarly, Rx2 can decode its desired signal concurrently by adjusting one of its antennas to align interference from Tx1 and Tx3; while, Rx3 adjusts one of its antennas to align interference from Tx1 and Tx2.

Reciprocity: We now have the receive antennas at positions where Tx1, Tx2 and Tx3 can transmit concurrently without interference. Suppose the receivers want to send ACKs back to their respective transmitters, concurrently. These ACKs may still interfere with each other at the transmitters. For example, Rx1’s ACK would cause interference at Tx2 and at Tx3. To address this issue, Rx1 can precode its transmission using standard interference nulling, such that it is nulled at Tx2. Fortunately, nulling the ACK at Tx2 also automatically nulls the ACK at Tx3 as well. This desirable phenomenon is due to channel reciprocity and will be further explained in §7. Hence, Rx2 and Rx3 can perform a similar procedure so that the receivers concurrently deliver ACKs to their transmitters, interference-free.

Multiplexing Gains: MoMIMO’s motion-based interference alignment can enable three concurrent streams in this setting, i.e., a $3\times$ gain over 802.11. We note that even if the receivers employed up-link multi-user MIMO (as in [26]) instead of standard 802.11, the maximum number of concurrent streams would be two, still $1.5\times$ less than that of MoMIMO.

Scaling: MoMIMO’s interference alignment scales beyond the setting in Fig. 4(b). For e.g., in a setup with four single-antenna transmitters and four 3-antenna receivers (shown in Fig. 4(c)), it can enable four concurrent streams, i.e., a $4\times$ gain over 802.11.

4.3 Combining Alignment and Nulling

MoMIMO’s throughput gains can be generalized to heterogeneous networks that require combining both motion-based alignment and motion-based nulling. Consider the scenario in Fig. 5, with two pairs of single-antenna nodes, and one pair of two-antenna MIMO nodes. Note that in 802.11n, only one of Tx1, Tx2, and Tx3 can transmit packets at any given time.

In contrast, MoMIMO can exploit both motion-based alignment and nulling to deliver three concurrent streams. To begin with, recall that Tx1, Tx2, Rx1 and Rx2 form a topology similar to the one in Fig. 4(a). Thus, to enable concurrent transmissions, Rx1 can simply slide its antenna to null Tx2, and Rx2 can similarly slide its antenna to null Tx1, as in §4.1.

We still need to make sure that Tx3 can transmit concurrently without interfering with the ongoing transmissions. To this end, we exploit interference alignment, as in §4.2. Specifically, Rx3 slides one of its antennas to align Tx1 and Tx2’s signals. In addition, Rx3




Figure 6: Interference Alignment. We wish to align the signals received from the two interferers I_1 and I_2 along the same direction. The goodness of alignment is determined by the *residual interference*; it is the projection of an interferer along a direction orthogonal to the desired alignment, and denoted as h_{orth} .

precodes its transmission using standard interference nulling, such that it is nulled at Tx1; by reciprocity, Rx3’s transmission is automatically nulled at Tx2 as explained in §7. Similarly, Tx3 slides one of its antennas to align the ACKs from Rx1 and Rx2; it also precodes its transmission such that it is nulled at Rx1 (and, hence, Rx2).

Multiplexing Gains: In this setting, 802.11n can send 1.33 concurrent streams on average assuming fair time-sharing between streams.⁵ In contrast, MoMIMO’s motion-based alignment and nulling enables three concurrent streams, i.e. a $2.25\times$ gain over 802.11n. Note that using interference alignment and nulling without motion in this setup will only enable two concurrent streams at any time. Hence, MoMIMO can still achieve a $1.5\times$ gain over advanced techniques that leverage standard nulling and alignment, e.g. n+ [15].

Finally, we note that in this section, we have assumed that it is feasible to find nearby locations that achieve the desired alignment and nulling. In §9, we revisit these topologies and provide extensive measurements that span line-of-sight and non-line of sight scenarios in multiple buildings, all of which confirm the gains of MoMIMO.

5. MOMIMO ANALYSIS

In this section, we extend signal propagation models to capture how multipath effects enable motion-based alignment. We also give insights for why the required displacement is relatively small and is about one or two inches.

In order to quantify alignment, we need to pick a metric that reflects the goodness of alignment. So, how do we choose this metric? Suppose, a receiver Rx wants to align a given interferer I along a particular direction \vec{v} in its antenna space. If the interferer is perfectly aligned with \vec{v} , then its projection along \vec{v}^\perp (i.e. the direction orthogonal to \vec{v}) is zero. In contrast, if the interferer is poorly aligned with \vec{v} , its projection on \vec{v}^\perp is large. In fact, the larger this projection is, the worse the alignment is. In the rest of this section, we analytically quantify the goodness of alignment in terms of *residual interference* after alignment [8, 20]. This quantity, which we denote h_{orth} is defined as the projection of the interference along the direction orthogonal to that of the desired alignment.

Consider the setup in Fig. 6, where we wish to slide antenna 1 of the receiver Rx to align the signal \vec{h}_1 from I_1 along the same direction as the signal \vec{h}_2 from I_2 . Let h_{ir} denote the channel from interferer I_i to receive antenna r . We then write the residual inter-

⁵In any time slot, Tx1, Tx2 or Tx3 sends 1, 1, or 2 streams respectively. Thus, on average, 802.11n sends $\frac{1+1+2}{3} = 1.33$ streams.




Figure 7: Specifying Reflectors. We specify the location of reflector i with respect to interferer I_i by: $\theta_{t1,i}$, the angle of departure of the signal from interferer I_i and $\phi_{t1,i}$, the angle of arrival of the reflected signal at receive antenna 1.

ference after alignment h_{orth} as:

$$h_{orth} = \vec{h}_1 \cdot \vec{h}_2^\perp \quad (1)$$

$$= (h_{11}, h_{12}) \cdot (-h_{22}, h_{21}) \quad (2)$$

$$= -h_{11}h_{22} + h_{12}h_{21} \quad (3)$$

where \vec{h}_2^\perp denotes the vector orthogonal to \vec{h}_2 .⁶

Let's study how the value of these channels h_r and the resulting h_{orth} change as a function of reflectors in the environment. Say that we have n reflectors in the environment. Then the channel between any interferer I_i and receive antenna r is the sum of the channel along the direct path between the two, $h_{r,0}$, plus the channels along all the reflected paths, $h_{r,i}$, for all reflectors $i = 1, \dots, n$. We can use standard propagation models [27] to express these channels as a function of path length. Specifically, the channel on the direct path $h_{r,0}$ is:

$$h_{r,0} = \frac{1}{d_{r,0}} e^{-2\pi j \frac{d_{r,0}f}{c}} \quad (4)$$

where $d_{r,0}$ is the length of the direct path between interferer I_i and receive antenna r , c is the speed of light, and f is the frequency of the transmitted signal.

Similarly, the channel along the i^{th} reflected path as:

$$h_{r,i} = \frac{\gamma_i}{d_{r,i}} e^{-2\pi j \frac{d_{r,i}f}{c} + j\pi} \quad (5)$$

where γ_i is the reflection coefficient of the reflector, and $d_{r,i}$ is the total distance traversed by the reflected signal – i.e., $d_{r,i}$ is the sum of the two segments of a reflected path, the segment from the interferer I_i to the reflector and the segment from the reflector to the receive antenna r .

Thus, the channel from interferer I_i to receive antenna r is:

$$\begin{aligned} h_r &= h_{r,0} + \sum_{i=1}^n h_{r,i} = \frac{1}{d_{r,0}} e^{-2\pi j \frac{d_{r,0}f}{c}} + \sum_{i=1}^n \frac{\gamma_i}{d_{r,i}} e^{-2\pi j \frac{d_{r,i}f}{c} + j\pi} \\ &= \sum_{i=0}^n \frac{\gamma_i}{d_{r,i}} e^{-2\pi j \frac{d_{r,i}f}{c} + \alpha_i} \end{aligned}$$

Where $\gamma_0 = 1$, $\alpha_0 = 0$, and $\alpha_i = j\pi$ for $i = 1, \dots, n$.

We can substitute these channels in Eqn. 3 that defined h_{orth} to understand how reflectors impact the residual interference after alignment. As we substitute, remember that we are only sliding antenna 1 on the receiver, while antenna 2 is fixed. Therefore, we consider channels h_{12} and h_{22} to antenna 2 to be constants, for the rest of our

⁶Note that, in practice, \vec{h}_2^\perp is a unit vector. We ignore normalization of \vec{h}_2^\perp for simplicity.

analysis.

$$\begin{aligned} h_{orth} &= -h_{22}h_{11} + h_{12}h_{21} \\ &= -h_{22} \sum_{i=0}^n \frac{\gamma_i}{d_{11,i}} e^{-2\pi j \frac{d_{11,i}f}{c} + \alpha_i} + h_{12} \sum_{i=0}^n \frac{\gamma_i}{d_{21,i}} e^{-2\pi j \frac{d_{21,i}f}{c} + \alpha_i} \\ &= \sum_{i=\{1,2\}} \sum_{i=0}^n \frac{\gamma_{t1,i}}{d_{t1,i}} e^{-2\pi j \frac{d_{t1,i}f}{c} + \alpha_i} \end{aligned}$$

Where $\gamma_{11,i} = -h_{22}\gamma_i$ and $\gamma_{21,i} = h_{12}\gamma_i$, are constants.

The above equation computes the residual interference after alignment for a particular position of the receive antenna. We are interested in how h_{orth} changes as the receive antenna slides by $(\delta x, \delta y)$, which we denote by $h_{orth}(\delta x, \delta y)$. To derive $h_{orth}(\delta x, \delta y)$, we express the location of reflector i based on two angles for any interferer I_i and receive antenna r : the angle of departure ($\theta_{tr,i}$) and angle of arrival ($\phi_{tr,i}$), as shown in Fig. 7. Given these values, we can show that $h_{orth}(\delta x, \delta y)$ is:

$$h_{orth}(\delta x, \delta y) = \sum_{i=\{1,2\}} \sum_{i=0}^n \frac{\gamma_{t1,i}}{d_{t1,i}} e^{-2\pi j \frac{d_{t1,i}f}{c} + \alpha_i} \quad (6)$$

Where $\tilde{d}_{t1,i}$ is given by:

$$\begin{aligned} \tilde{d}_{t1,i} &= \sqrt{(2c_{t1,i}\sin^2\beta_{t1,i} - d_{t1,0} - \delta x)^2 + (c_{t1,i}\sin(2\beta_{t1,i}) - \delta y)^2} \\ c_{t1,i} &= \frac{d_{t1,0}\sin\phi_{t1,i}}{\sin\phi_{t1,i} - \sin\theta_{t1,i}} \\ \beta_{t1,i} &= (\phi_{t1,i} - \theta_{t1,i})/2 \end{aligned}$$

The proof for the above equations follows from the geometry of the reflectors in Fig. 7. It is provided in Appendix A.

Eqn. 6 shows how sliding a receive antenna impacts alignment, or more precisely, the residual interference orthogonal to the desired alignment. As Eqn. 6 is fairly complex, in the following sections, we extract insights from it considering specific scenarios.

5.1 Impact of a Reflector on Motion-Based Alignment

We first consider the case of single randomly positioned non-absorbent reflector (i.e., $\gamma_i = 1$). Of course in reality, an indoor scenario would have many absorbent reflectors. However, studying this simplified case allows us to understand the trends that control how a reflector impacts alignment. Later in §5.2, we use these insights to understand the impact of many absorbent reflectors.

Recall that our objective is to understand, given that the receive antenna is at an arbitrary initial position, how much displacement is needed to move it to a location that achieves alignment (i.e., minimizes $|h_{orth}|^2$, the residual interference after alignment). To compute this displacement, we numerically compute $|h_{orth}|^2$ using Eqn. 6. Recall that the impact of the reflector on the signal from each interferer I_i is defined by the angle of arrival (ϕ_{t1}) and the angle of departure (θ_{t1}) of the signal, as shown in Fig. 7. Since we have two interferers and one reflector, in this example, we have four parameters: θ_{11} , θ_{21} , ϕ_{11} and ϕ_{21} , which determine this displacement.

Fig. 8(a) plots the mean displacement from an arbitrary initial position of the receive antenna to a position that achieves alignment by minimizing $|h_{orth}|^2$ locally.⁷ The figure shows that the displacement is small when either ϕ_{11} or ϕ_{21} is large. Specifically, when either ϕ_{11} or ϕ_{21} is between 90° and 180° , the required mean displacement varies between about 0.6 inches and 1.2 inches. We note

⁷Our analysis of h_{orth} also reveals that these points of local minima, on average, decrease $|h_{orth}|^2$ below the noise floor, provided the environment has reflectors close to the receiver, that are behind it with respect to either sender.

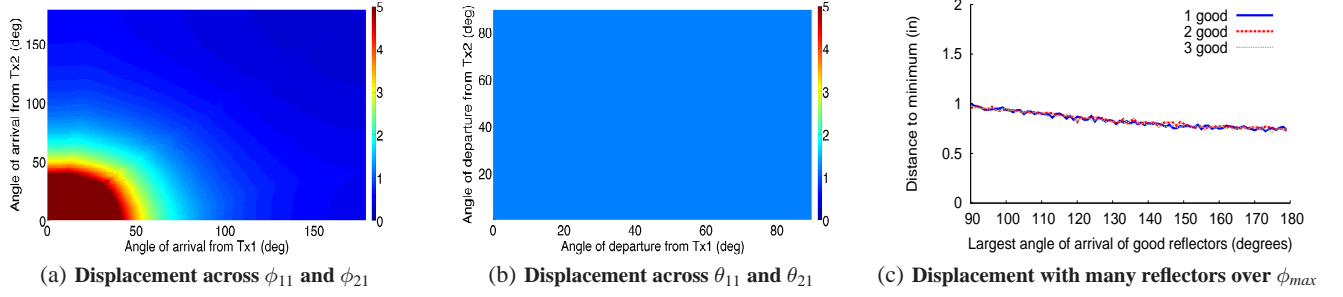


Figure 8: Displacement (in inches) to position of alignment (a): Across ϕ_{11} and ϕ_{21} , with fixed arbitrary θ_{11} and θ_{21} . We cap the maximum at 5 inches for ease of viewing. (b) Across θ_{11} and θ_{21} , with $\phi_{11}, \phi_{21} = 90^\circ$. Geometry restricts $\theta_{11}, \theta_{21} \leq 90^\circ$. (c) With 10 reflectors placed at random locations across ϕ_{max} , the largest angle of arrival among the good reflectors with respect to any interferer.

that this displacement and the results in Fig. 8(a) are based on Eqn. 6 for $f = 2.5$ GHz, which is the typical carrier frequency of 802.11. Hence the displacements of 0.6 and 1.2 inches correspond to distances of $\lambda/8$ and $\lambda/4$ respectively. The required displacements are even smaller at higher frequencies, where the wavelength is shorter, for e.g., in the 5 GHz Wi-Fi band.

So what does ϕ_{11} or $\phi_{21} > 90^\circ$, mean in terms of the physical location of the reflector? Recall from Fig. 7 that ϕ_{i1} denotes the angle along which reflected signal arrives at the receiver from interferer I_i . Thus, if $\phi_{i1} > 90^\circ$, then this reflector is behind the receiver with respect to interferer I_i . Hence, we arrive at the following insight:

Insight 1: Reflectors which are behind the receiver with respect to the interferer (i.e. ϕ_{11} or $\phi_{21} > 90^\circ$) enable a displacement on the order of 1 or 2 inches to a position that enables alignment. We call such reflectors *good reflectors*.

Note that it is common for most receivers in wireless networks to be placed on some platform (e.g. a table, the floor, etc.), which is usually behind the receiver with respect to at least one interferer, and hence serves as a good reflector.

One may wonder if the displacement to a position of alignment depends on either θ_{11} or θ_{21} , given that ϕ_{11} or ϕ_{21} are fixed at values $\geq 90^\circ$. In fact, we observe that the angles of departure have no impact on the displacement, once we fix the angles of arrival. In particular, Fig. 8(b), demonstrates that the required mean displacement is constant versus θ_{11}, θ_{21} , for the case when ϕ_{11} and ϕ_{21} are both fixed at 90° .

5.2 Motion-Based Alignment with Multiple Reflectors

While our discussion so far focused on a single reflector, the natural question to ask is how does the displacement to a position of alignment change, given a large number of reflectors in the environment. To this end, we extend our analysis to multiple absorbent reflectors. In particular, we pick multiple reflectors with random angles of departure and arrival and an absorption coefficient chosen randomly between $[0.2, 0.6]$ [14]. We use our model in Eqn. 6 to evaluate how these reflectors affect the mean displacement required from an arbitrary location to reach a position that achieves alignment by locally minimizing $|h_{orth}|^2$.

We consider a scenario with 10 randomly positioned reflectors with different angles of arrival and departure. We consider the following three cases: (1) Only one of these reflectors is a “good reflector” with respect to either interferer I_1 or interferer I_2 . (2) Two reflectors are “good reflectors” with respect to either interferer. (3) Three reflectors are “good reflectors” with respect to either interferer.

Fig. 8(c) plots the mean displacement needed to reach a position of alignment, i.e. minimize $|h_{orth}|^2$. We plot this quantity across the largest angle of arrival, $\phi_{max} = \max_{i \in \{1,2\}, i} \phi_{i1,i}$, among all the good

reflectors with respect to either interferer. The figure reveals the following insight:

Insight 2: As long as at least one good reflector exists in the environment, i.e. one reflector is behind the receiver with respect to either interferer, the mean displacement needed to reach a position of alignment, is in the vicinity of $\lambda/8$ to $\lambda/4$, i.e. around 1 or 2 inches, for 2.4-2.5GHz transmissions.

While our analysis so far has focused on interference alignment, it can readily be extended to interference nulling. In particular, we can consider nulling as a special case of alignment where h_{orth} is the projection along the direction of the receive antenna where the signal needs to be nulled (i.e., one of the axes in the antenna space). Hence, the analysis of h_{orth} is similar.

6. STOCHASTIC HILL CLIMBING

In this section, we present a Stochastic Hill Climbing algorithm that automatically adjusts the receive antenna to seek positions of interference alignment or nulling. To find these positions, the algorithm minimizes the *Interference to Noise Ratio* (INR) from the interfering transmitter(s). For interference nulling, INR is the ratio of the power of the interfering transmitter to the power of noise. For interference alignment, we define the INR as follows. Let (α, β) denote the direction along which the projected power from the interfering signals is minimum. Recall that the received signal must be projected along this direction, to decode the desired signal. Thus, we define the INR for alignment as the ratio of the power of the interfering signals projected along this direction, to the noise power σ^2 . For a system with two single-antenna transmitters and a two-antenna receiver, the INR can be computed based on the wireless channels as:

$$INR = \frac{1}{\sigma^2} \min_{\alpha, \beta} |h_{11}\alpha + h_{12}\beta|^2 + |h_{21}\alpha + h_{22}\beta|^2$$

where h_{ij} is the channel from transmitter i to receive antenna j in Fig. 6. The solution of this minimization problem has a closed form (Proof in Appendix B):

$$INR = \frac{1}{2\sigma^2} \left(|h_{11}|^2 + |h_{21}|^2 + |h_{12}|^2 + |h_{22}|^2 \right) - \frac{1}{\sigma^2} \sqrt{\left| \frac{|h_{11}|^2 + |h_{21}|^2 - |h_{12}|^2 - |h_{22}|^2}{2} \right|^2 + |h_{11}h_{12}^* + h_{21}h_{22}^*|^2} \quad (7)$$

Because channels are continuous functions over space, the INR profile is continuous and smooth. In other words, an incremental movement in any given direction would lead to a gradual increase or decrease in INR. Hence, to minimize the INR, a receiver can slide its antenna gradually until it reaches a local minimum.

We leverage this intuition in designing our stochastic hill climbing algorithm to reduce the INR (pseudo-code in Alg. 1). At a high level, our algorithm proceeds as follows: the receiver continuously monitors the INR from its interferer(s) (using Eqn. 7). Initially, it picks an arbitrary direction, and slides one of its antennas in that direction. Specifically, it slides the antenna either forward or backward, depending on which of those decreases the INR. It continues to move the antenna along that direction until the INR reaches a local minimum (in that direction). Then, it rotates the antenna either clock-wise or counter clock-wise about a small arc (of radius a few mm), whichever decreases the INR. Again, the receiver keeps rotating that antenna until the INR reaches a local minimum. The receiver repeats the above process, switching between translations and rotations, until it reaches a local minimum (in both translation and rotation). If the INR at the local minimum is below the noise floor, the algorithm terminates. Otherwise, the receiver moves the antenna to a random location and repeats the process again.

In practice, across all experiments in §9, the stochastic hill climbing algorithm converges to an INR below the noise floor in an average of 3 tries.

Algorithm 1 Pseudo-code for Stochastic Hill Climbing

```

mode = +1                                ▷ 1: Translate, -1: Rotate
direction = +1                            ▷ 1: Forward, -1: Reverse
attempts = 0
▷ Loop while INR is above 0 dB
while INR > 0 dB do
  switch mode
    case +1: TRANSLATE(direction);        ▷ In Translate Mode
    case -1: ROTATE(direction);          ▷ In Rotate Mode
  end switch
  prev_INR = INR                          ▷ Previous INR of sender
  INR = GETINR()                          ▷ Update INR based on channels
  if INR > prev_INR then                  ▷ If INR increased
    direction = -direction                ▷ Switch direction
    attempts = attempts + 1
  end if
  if attempts == 2 then                  ▷ Hit a minimum
    mode = -mode                          ▷ Switch modes
    attempts = 0
  end if
  if At local minimum with INR > 0 dB then
    ▷ Move to Random Location
  end if
end while

```

7. RECIPROCIDY


In this section, we show how MoMIMO leverages channel reciprocity to allow concurrent uplink transmissions (such as ACKs). Consider two wireless nodes i and j . Let h_{ij} and h_{ij}^{rev} denote the forward channel from i to j and reverse channel from j to i , respectively. Then reciprocity [7] states that:

$$h_{ij}^{rev} = t_i h_{ij} r_j \quad (8)$$

Where t_i and r_j are fixed constants that depend purely on nodes i and j respectively.

Eqn. 8 has an interesting consequence for interference nulling. Specifically, when MoMIMO adjusts j 's receive antenna so that $h_{ij} \approx 0$, this also results in $h_{ij}^{rev} \approx 0$. Hence, nulling the downlink channel from interferer i to receiver j by motion, also nulls the uplink channel from j to i .

Reciprocity can also be exploited for interference alignment. Consider two single-antenna nodes I_1 and I_2 causing interference at a two-antenna receiver Rx, as shown in Fig. 6. MoMIMO moves one of Rx's receive antennas to a position of interference alignment,



(a) Testbed in Bldg. 1 (b) Testbed in Bldg. 2

Figure 9: Testbed topology. The figure depicts the cropped floor plans of two buildings where the experiments were conducted. The red marks on the floor plans depict locations for wireless nodes. Different random subsets of these locations for the transmitters and receivers are picked for different experiments.

so that (Eqn. 3):

$$h_{11} h_{22} - h_{21} h_{12} \approx 0 \quad (9)$$

Now, suppose the two-antenna receiver Rx needs to send a single stream on the uplink without interfering with I_1 or I_2 . It precodes its signal by $(-h_{12}^{rev}, h_{11}^{rev})$ so that the effective uplink channel at interferer I_1 : $-h_{12}^{rev} h_{11}^{rev} + h_{11}^{rev} h_{12}^{rev} = 0$, is nulled. Interestingly, because the receive antenna is in a position of alignment at the downlink, this precoding also nulls the effective uplink channel at I_2 . Specifically, from Eqn. 8 and 9, the effective uplink channel at I_2 is:

$$\begin{aligned}
-h_{12}^{rev} h_{21}^{rev} + h_{11}^{rev} h_{22}^{rev} &= -t_1 h_{12} r_2 t_2 h_{21} r_1 + t_1 h_{11} r_1 t_2 h_{22} r_2 \\
&= t_1 t_2 r_1 r_2 (h_{11} h_{22} - h_{21} h_{12}) \approx 0
\end{aligned}$$

Thus, MoMIMO's interference alignment and nulling by motion provides not only concurrent downlink transmissions, but by leveraging precoding and reciprocity, also enables concurrent uplink transmissions.

8. EXPERIMENTAL SETTING

We describe our experimental environment, which is used to illustrate the gains of MoMIMO.

Implementation: Nodes in our experiments are equipped with USRP-N210 software radios [5] and RFX 2400 daughter-boards. They communicate in the 2.4 GHz Wi-Fi band using 20 MHz signals. We implement OFDM directly in the USRP Hardware Driver (UHD). We use various 802.11 modulations (BPSK, 4QAM, 16QAM, and 64QAM), coding rates, and choose between them using the effective-SNR bitrate selection algorithm [10]. Our implementation periodically measures the channels by listening to packets from the desired and interfering transmitters. It tracks the INR of the interfering signals for interference alignment or nulling.

To emulate sliding antennas, the antenna of each USRP is mounted on an iRobot Create robot, controlled by an ASUS EEPIC 1015PX netbook. We implement stochastic hill climbing in a fully distributed manner using iRobot Create's Open Interface. Each receiver measures INR for interference alignment or nulling in real-time from the software radio(s), and automatically maneuvers the robot to minimize INR from any interferer(s). It also queries the robot sensors to ensure that the robot does not move beyond 2 inches from its initial position as it seeks to minimize its INR.

Tested Environments: We evaluate MoMIMO in indoor settings, in line-of-sight (LOS) and non-line-of-sight (NLOS) scenarios, and in two different buildings.

Our main experiments were conducted in Building 1 depicted in Fig. 9(a), which hosts a computer science department. The building is built mostly of flat reinforced concrete slabs and columns.




Figure 10: INR before and after applying MoMIMO’s alignment or nulling. The figure depicts the CDF of motion-based alignment in (a), (b) and nulling in (c), (d). Our experiments are performed in line-of-sight scenarios in (a), (c) and non-line-of-sight scenarios in (b), (d).

The enclosure of the exterior walls is of metal and brick. The interior walls are gypsum with steel sheet metal studs. The rooms are standard offices with several desks, chairs, cabinets, and computer workstations.

We also ran experiments in Building 2, which is constructed in 1963, and has thick concrete interior walls. We performed our experiments in a large classroom connected to a long corridor with the layout as shown in Fig. 9(b). The classroom consists of several desks and chairs.

Metrics: To evaluate MoMIMO’s effectiveness at achieving interference alignment and nulling, we measure the *interference to noise ratio (INR)*. In the case of nulling, the INR is computed as the total received interference power after running the stochastic hill climbing algorithm. In the case of alignment, the INR is measured after running the stochastic hill climbing algorithm by projecting the interference signals orthogonal to the desired alignment direction, and computing the interference power after projection. In both cases, an INR less than 0 dB means that MoMIMO has reduced the interference below the noise level.⁸

Other metrics of interest include the throughput, which we measure using the effective-SNR metric as described in [10], and the required antenna displacement for nulling or alignment, which we measure as the distance from the initial location of the antenna to its location after running the stochastic hill climbing algorithm.

9. EMPIRICAL RESULTS

We evaluate the performance of MoMIMO in the testbed environment described in §8.

9.1 Microbenchmarks

We first investigate whether local antenna adjustments, using our stochastic hill climbing algorithm, can indeed deliver interference alignment and nulling. To do so, we quantify the INR after MoMIMO’s alignment and nulling, and check whether it can be reduced below the noise level, i.e., 0 dB.

Interference Alignment: We experiment with the setup in Fig. 6, which consists of two single-antenna interferers and a 2-antenna receiver. We place these nodes randomly in our testbeds in both line-of-sight (LOS) and non-line of-sight (NLOS) scenarios. The two interferers concurrently transmit packets back-to-back. The receiver estimates the total power received on both of its antennas; this power constitutes the initial INR. The receiver then runs stochastic hill climbing until the two received signals are aligned. Subsequently, it computes the final INR as the projection of the

⁸To be able to evaluate the INR accurately, packets sent by the interferer contain a header with twenty known OFDM symbols. This allows us to measure INR up to an accuracy of -13 dB by averaging the wireless channel across these repeated symbols. While this allows us to accurately measure the residual interference, in practice, pushing the INR below -3 dB is unnecessary. Thus 2-3 known OFDM symbols are sufficient for this accuracy.

received signals along the direction orthogonal to the desired alignment as described in §6. We repeat this experiment 20 times, each time placing the interferers and receiver at different locations in our topology.

Figs. 10(a) and 10(b) show the cumulative distribution functions (CDFs) of both the initial and final INRs in these experiments for LOS and NLOS scenarios. The figures show an average INR reduction of about 22 dB in both LOS and NLOS scenarios. Note that the median residual INR is -0.5 dB and -2.5 dB, in LOS and NLOS respectively, which is below the noise floor, thereby enabling significant throughput gains.

The CDFs also show that in a few percent of the runs, the INR after alignment may be above the noise level by a few dBs. These runs correspond to cases of very strong interference, over 25 dB, and hence the few dB of residual INR. Even in these scenarios, the reduction in INR is large and exceeds 20 dB. Overall, the figure shows that motion-based alignment produces similar accuracy to multi-antenna alignment [15], but requires only single antenna transmitters.

Interference Nulling: We place a single-antenna interferer and a single-antenna receiver randomly in our testbed, in both LOS and NLOS scenarios. The interferer transmits a stream of packets back-to-back. The receiver slides its antennas following the stochastic hill climbing algorithm in order to null the signal from the interferer. Upon convergence, it computes the final INR. We repeat this experiment 20 times, each time placing the nodes in different locations in our testbed.

Figs. 10(c) and 10(d) show the CDFs of the initial and final INRs in these experiments for both line-of-sight (LOS) and non-line-of-sight (NLOS) scenarios. The figure shows that after MoMIMO nulling, the median residual INR is -0.3 dB in LOS and NLOS, and median reduction in INR is about 15 dB. Our results show that motion-based nulling effectively cancels interference and brings it to the noise level.

In comparison with the alignment experiments, we note that the initial INR for nulling is much lower. This is because in alignment, the initial INR is due to two interferers, whereas here it is due to only one interferer.

9.2 Throughput Comparison

We assess the throughput benefits of MoMIMO against two baselines: standard 802.11n and n+ [15]. We compare against n+ to show that MoMIMO delivers gains even when compared with a system that employs advanced MIMO techniques.

For this experiment, we use the heterogeneous network topology in Fig. 5. We note that other topologies like those in Fig.4(b) and 4(c) show a bigger gain for MoMIMO, as explained in §4. However, we chose to present empirical results for the heterogeneous network topology because it is more complex than the others and requires combining both motion-based nulling and alignment.

Recall from §4 that the heterogeneous network in Fig. 5 has two single-antenna transmitter-receiver pairs and one 2-antenna




Figure 12: Leveraging Reciprocity. CDF of INR and Reciprocal INR for interference alignment and nulling.




Figure 13: Displacement for INR reduction. CDF of the total displacement of the receiver from its initial position, across multiple experiments, for interference alignment or nulling.




Figure 14: INR across Bandwidth. Plot of the mean and standard deviation of the INR at the converged position for interference alignment across bandwidth.




Figure 11: Total Network Throughput. Total network throughput of all three schemes: 802.11n, n+, and MoMIMO.

transmitter-receiver pair. For this network, MoMIMO uses sliding antennas to null the signal of Tx1 at Rx2 and that of Tx2 at Rx1. It also needs to align the signals of Tx1 and Tx2 at Rx3. In this case, MoMIMO can deliver 3 concurrent streams. In contrast, 802.11n has to budget the wireless medium between the three Tx-Rx pairs which yields an average of 1.33 concurrent streams, while n+[15] can enable 2 concurrent streams.

For each run of our experiments we pick a random subset of the nodes in the testbed to represent the nodes in the heterogeneous network in Fig. 5. We consider a single long-lasting flow from each sender to its receiver. We repeat the experiment at 20 randomly chosen locations. At each location, we record throughputs for MoMIMO, 802.11n and n+.

For each of the MoMIMO runs, the network initially uses 802.11n with each pair transmitting in a different time slot. During this phase, the nodes do not transmit concurrently; instead they measure the channel and run stochastic hill climbing. Once stochastic hill climbing reaches the desired nulling and alignment, we allow the nodes to transmit concurrently.

Fig. 11 shows CDFs of the obtained throughput in MoMIMO as well as those of 802.11n and n+. On average, MoMIMO achieves a throughput gain of $1.98\times$ over 802.11 and $1.31\times$ gain over n+. In comparison, the expected throughput gains for this heterogeneous network as reported in §4.3 are $2.25\times$ over 802.11n and $1.5\times$ over n+. Thus, our experimental gains are close to the expected ones. Note that both MoMIMO and n+ suffer from a slight performance dip in comparison to their theoretical gains; this is due to small residual interference power after alignment and nulling.

9.3 Reciprocity

Reciprocity enables MoMIMO to perform motion-based alignment/nulling on the forward path and still enjoy alignment and nulling on the reverse path, thereby enabling ACKs to be transmitted concurrently without interference. In this section, we evaluate how eliminating interference by applying MoMIMO on the forward path translates into eliminating interference on the reverse path.

Here, we use the same experiments described in §9.1 for both alignment and nulling. The difference however is that after we align/null along the forward path, we reverse the direction of the transmission, i.e., make the receiver transmit and the previous transmitter receive. For the case of alignment and as described in §7, the receiver also precodes the ACK to null it at one of the aligned senders. Our objective is to show empirically that this naturally leads to nulling the signal at the other sender. Hence we measure the corresponding reciprocal INR. For nulling we simply want to show that nulling in the forward direction leads to nulling in the reverse direction. Hence we measure the reverse INR.

Fig. 12 plots the CDFs of the reverse INRs for both alignment and nulling. Each CDF is taken over 20 different runs with different node locations. The figure shows that indeed applying stochastic hill climbing in the forward direction virtually eliminated the INR in the reverse direction as well, reducing it below the noise level (i.e., below 0 dB).

9.4 Displacement

Next, we would like to quantify the amount of displacement by which MoMIMO needs to slide the antennas to achieve alignment or nulling. To do so, we measure the displacement between the initial and final positions of the receive antennas in the throughput experiments in §9.2.

Fig. 13 plots the CDFs of this displacement for both alignment and nulling. As expected, since we restricted the motion of the antennas to within a 2 inch radius, the maximum displacement across all experiments is 2 inches. Further, the figure shows that MoMIMO's stochastic hill climbing algorithm needs a mean displacement of 0.44 inches for interference alignment and 1.17 inches for interference nulling. Alignment is easier to satisfy as many options exist for the direction along which two senders may be aligned, which provides extra flexibility for the choice of channels.

9.5 Impact of Channel's Bandwidth

Our previous experiments use a bandwidth of 20 MHz, the common setting for 802.11. In this experiment, we evaluate how channel bandwidth impacts MoMIMO's performance.

As in §9.1, we consider two single-antenna interferers and a two-antenna receiver. The receiver performs stochastic hill climbing by moving one of its antennas to align the signals from the two single-antenna interferers. Upon convergence, we fix the initial position of the receiver and repeat the experiment for different channel bandwidths: 2MHz, 4MHz, 8MHz, 16MHz, and 20MHz. For each channel width, we repeat this experiment in 10 randomly chosen locations in our testbed.

Fig. 14 plots the mean and standard deviation of the INR at the receiver as a function of the bandwidth, averaged over all our experiments. The plot reveals that the INR rises by an average of

5 dB as the bandwidth is increased from 2MHz to 20MHz. This is expected because a receiver perceives greater frequency diversity across OFDM subcarriers with increased bandwidth. Nevertheless, even at 20MHz, the mean INR across these experiments was around -2 dB, which is well below the noise floor. Thus, even though the performance of MoMIMO is impacted by an increase in bandwidth due to frequency diversity, MoMIMO continues to provide a significant reduction in INR.

9.6 Performance in Dynamic Environments

In this experiment, we study how MoMIMO responds to changes in the wireless channel. MoMIMO nodes incorporate a fail-safe mode that responds to channel changes by defaulting to 802.11n and re-running the stochastic hill climbing algorithm. Specifically, during the periods in which the receive antenna is still moving and has not reached a nulling/alignment position, as well as in which the nulling/alignment position has changed, the nodes fall back to 802.11n and do not transmit extra concurrent streams beyond what is allowed by 802.11n. Note that this behavior is natural in MoMIMO since, once the nulling does not hold, the receiver immediately starts hearing the interference. Similarly, if the alignment does not hold, the receiver sees the interference signal in the whole 2-D space as opposed to along a single direction. Thus, if the channel changes enough to disturb alignment or nulling, the receiver can detect this effect in real-time and fall back to standard 802.11n.

To assess MoMIMO’s fail safe mechanism, we repeat the experiment in §9.2, which uses the heterogeneous network. We focus on the two-antenna receiver (Rx3 from Fig. 5), which slides its antenna to align the signals from the two single-antenna interferers Tx1 and Tx2. We consider a scenario where the interferers (Tx1 and Tx2) are inside an office with an open door and the receiver is outside the office. Before starting the experiment, the receiver runs the stochastic hill climbing algorithm to align the two single-antenna transmitters. Our experiment spans 30 seconds. At $t = 15.4$ seconds, we close the door so that the interferers (Tx1 and Tx2), and the receiver are no longer in line-of-sight of each other. We track the INR of alignment at the receiver for the full duration of 30 seconds. We also measure the receiver’s throughput obtained based on the effective SNR [10] metric.

Fig. 15(a) plots a trace of the measured INR over time. The trace depicts a distinct spike of about 8 dB in INR at $t = 15.4$ seconds, when the door is closed. However the stochastic hill climbing algorithm responds quickly, reducing the INR to around 0 dB by $t = 16.6$ seconds. Note that the delay of 1.2 seconds was partly due to the limitations of the iRobot Create which is programmed to move at an average speed of 20 mm/s for accurate maneuverability.

Fig. 15(b) plots the throughput of the MIMO receiver over time, under identical settings for 802.11 and MoMIMO. At $t = 15.4$ seconds, the spike in INR caused by the change in channel triggers MoMIMO’s fail-safe mechanism. Hence, MoMIMO ceases concurrent transmissions, and falls back to 802.11n’s throughput levels (around 17 Mb/s on average). During this period, stochastic hill climbing moves the antenna but without activating concurrent transmissions. Once stochastic hill climbing causes the INR to fall close to the noise level MoMIMO switches on concurrent transmissions. By $t = 16.6$ seconds, the average throughput is around 24 Mb/s, restoring MoMIMO’s original gains over 802.11n for the new channel.

10. RELATED WORK

Related work falls in the following two categories:

(a) **MIMO Interference Management:** Recent years have witnessed advances in MIMO interference management from both the

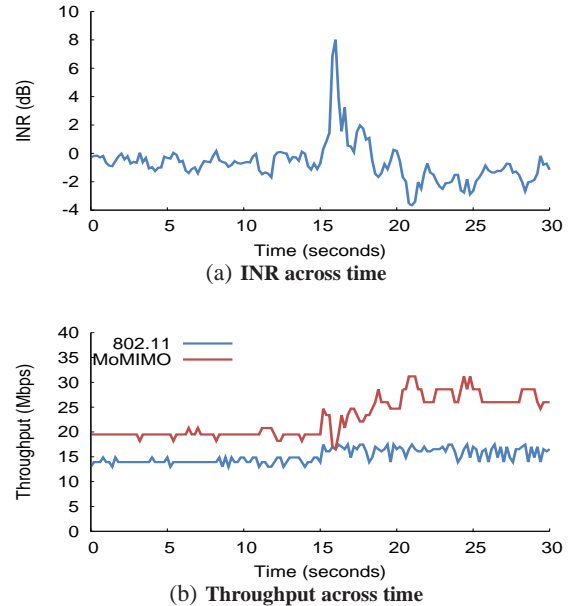


Figure 15: Performance of Interference alignment in Dynamic Environments. The plot shows a trace of the INR and throughput for 30 seconds, where at $t = 15.4$ s, the channels change from line-of-sight to non-line-of-sight.

theoretical [17, 2] and the empirical research communities [26, 7, 15, 22, 6, 21, 23]. Past systems implementing interference alignment [15, 7, 13] typically require multi-antenna senders that are aware of channel state information. Blind interference alignment [18], does not require channel state information at the transmitter, but needs symbol-level time synchronization between transmitters and receivers. Some theoretical work proposes interference alignment for single-antenna nodes either in frequency [25], or in time [19]. In contrast, MoMIMO is the first system to demonstrate that interference alignment and nulling can be performed purely by moving the receive antennas by just about one or two inches.

(b) **Exploiting Motion to Improve Wireless Performance:** Prior work demonstrated that mobility improves the performance of wireless network. For instance, Grossglauser and Tse [9] show that a wireless network of freely moving nodes has a higher capacity than a network with stationary nodes. Also, research from the robotics community [1, 12, 16] has leveraged mobility to improve the quality of wireless channels. For example, in [1, 12], robots improve their connectivity by maintaining a line-of-sight path to their access points. Other work observes that a robot can improve its throughput by sampling different locations in the environment, and choosing the one that maximizes its throughput (due to a higher SNR) [16]. Furthermore, recent work has leveraged node movement to improve beamforming [4] and energy consumption [29]. Additionally, some recent products allow the user to manually slide the antenna to avoid dead spots and improve the SNR [28].

MoMIMO builds on this past work but differs from it by being the first to show that motion enables MIMO-type interference alignment and nulling. Said differently, past work uses motion to improve the SNR but does not show that motion enables MIMO multiplexing gains.

11. CONCLUSION

In this paper, we introduced MoMIMO, a technique that demonstrated, for the first time, that interference alignment and nulling can be achieved, even with single-antenna transmitters, by simply

moving the receive antenna. We also showed that the amount of antenna displacement needed is fairly small (\sim one inch); hence, MoMIMO can be achieved by sliding antennas.

This paper focused on demonstrating the feasibility and potential gains of MoMIMO. Important topics for future work include: 1) designing a medium access protocol that leverages the synergy between motion and interference alignment as demonstrated by MoMIMO; and 2) studying MoMIMO's gains in different classes of dynamic environments, e.g. static vs. mobile clients.

We believe MoMIMO presents a new paradigm for interference management, especially in settings of long-lived interference patterns such as home networks and data-center flyways. Further, we envision that future systems may exploit motion-based interference alignment and nulling to enable new applications at the intersection of networking and robotics, where mobility is already innate.

12. ACKNOWLEDGEMENTS

We thank Omid Abari, Haitham Hassanieh, Ezz Hamad, Jue Wang, Arthur Berger, Diego Cifuentes, Peter Iannucci, Zack Kabelac, Nate Kushman, Jonathan Perry, Hariharan Rahul, Lixin Shi, the reviewers, and our shepherd, Li Erran Li, for their insightful comments. This research is supported by NSF. We thank members of the MIT Center for Wireless Networks and Mobile Computing: Amazon, Cisco, Google, Intel, Mediatek, Microsoft, ST Microelectronics, and Telefonica for their support.

13. REFERENCES

- [1] R. Arkin and T. Balch. Line-of-sight constrained exploration for reactive multiagent robotic team. *AMC*, July 2002.
- [2] V. Cadambe and S. Jafar. Interference alignment and spatial degrees of freedom for the k user interference channel. In *IEEE ICC*, 2008.
- [3] Samsung wireless smartcam. <http://www.samsungsv.com/Model/Detail/26/SNH-1011N>.
- [4] N. Chatzipanagiotis, Y. Liu, A. Petropulu, and M. M. Zavlanos. Controlling groups of mobile beamformers. In *IEEE CDC*, 2012.
- [5] USRP N210. <http://www.ettus.com>. Ettus Inc.
- [6] S. Gollakota, F. Adib, D. Katabi, and S. Seshan. Clearing the RF smog: making 802.11n robust to cross-technology interference. In *ACM SIGCOMM*, 2011.
- [7] S. Gollakota, S. Perli, and D. Katabi. Interference Alignment and Cancellation. In *ACM SIGCOMM*, 2009.
- [8] K. Gomadam, V. R. Cadambe, and S. A. Jafar. Approaching the capacity of wireless networks through distributed interference alignment. In *IEEE GLOBECOM*, 2008.
- [9] M. Grossglauser and D. N. C. Tse. Mobility increase the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking*, 2002.
- [10] D. Halperin, W. Hu, A. Sheth, and D. Wetherall. Predictable 802.11 packet delivery from wireless channel measurements. In *ACM SIGCOMM*, 2010.
- [11] D. Halperin, S. Kandula, J. Padhye, P. Bahl, and D. Wetherall. Augmenting data center networks with multi-gigabit wireless links. In *ACM SIGCOMM*, 2011.
- [12] M. A. Hsieh, A. Cowley, V. Kumar, and C. J. Taylor. Towards the deployment of a mobile robot network with end-to-end performance guarantees. 2006.
- [13] S. Kumar, D. Cifuentes, S. Gollakota, and D. Katabi. Bringing cross-layer MIMO to today's wireless LANs. In *ACM SIGCOMM*, 2013.
- [14] O. Landron, M. Feuerstein, and T. Rappaport. A comparison of theoretical and empirical reflection coefficients for typical exterior wall surfaces in a mobile radio environment. *IEEE Transactions on Antennas and Propagation*, 1996.
- [15] K. C.-J. Lin, S. Gollakota, and D. Katabi. Random access heterogeneous MIMO networks. In *ACM SIGCOMM*, 2011.
- [16] M. Lindhe, K. H. Johansson, and A. Bicchi. An experimental study of exploiting multipath fading for robot communications. In *Robotics: Science and Systems*, 2007.
- [17] M. Maddah-Ali, A. Motahari, and A. Khandani. Communication over MIMO X channels: Interference alignment, decomposition, and performance analysis. *Information Theory, IEEE Transactions on*, 2008.
- [18] K. Miller, A. Sanne, K. Srinivasan, and S. Vishwanath. Enabling real-time interference alignment: promises and challenges. *MobiHoc*, 2012.
- [19] 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. *CoRR*, 2009.
- [20] S. W. Peters and R. W. Heath. Cooperative algorithms for MIMO interference channels. *IEEE Transactions on Vehicular Technology*, 2011.
- [21] H. Rahul, S. S. Kumar, and D. Katabi. Jmb: Scaling wireless capacity with user demand. In *ACM SIGCOMM*, 2012.
- [22] W.-L. Shen, Y.-C. Tung, K.-C. Lee, K. C.-J. Lin, S. Gollakota, D. Katabi, and M.-S. Chen. Rate adaptation for 802.11 multiuser MIMO networks. In *ACM MobiCom*, 2012.
- [23] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong. Argos: Practical many-antenna base stations. In *ACM MobiCom*, 2012.
- [24] Sony bravia wireless link module. <http://www.lg.com/us/tv-accessories/lg-AN-WL100W-wireless-media-kit>.
- [25] C. Suh and D. Tse. Interference alignment for cellular networks. In *Allerton*, 2008.
- [26] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. M. Voelker. SAM: enabling practical spatial multiple access in wireless LAN. In *ACM MobiCom*, 2009.
- [27] D. Tse and P. Vishwanath. *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
- [28] Netgear a6200 usb wi-fi adapter with sliding antennas. www.netgear.com/home/products/wireless-adapters/ultimate-wireless-adapters/a6200.aspx.
- [29] Y. Yan and Y. Mostofi. Co-optimization of communication and motion planning of a robotic operation in fading environments. In *IEEE ASILOMAR*, 2011.

APPENDIX

A. LENGTH OF REFLECTED PATHS

In this section, we derive the length of reflected paths given the position of an arbitrary reflector. Recall from §5 that any reflector i can be specified in terms of two angles $\theta_{r1,i}$ and $\phi_{r1,i}$. Consider a transmitter I_t and receive antenna R_1 separated by a distance $d_{t1,0}$. We can write the position of the reflector in terms of $c_{r1,i}$, its distance from the transmitter, and $\beta_{r1,i}$ its angle to the axis I_tR_1 (Figure 16). First, it is easy to see that $\beta_{r1,i} = \frac{\phi_{r1,i} - \theta_{r1,i}}{2}$. Now, by using sine-rule on $\triangle PI_tR_1$, we have:

$$PI_t = d_{t1,0} \frac{\sin(\phi_{r1,i})}{\sin(\phi_{r1,i} + \theta_{r1,i})} \quad (10)$$

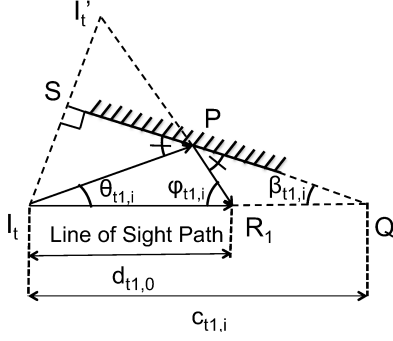


Figure 16: Reflected Path. We consider a transmitter I_t and receive antenna R_1 . Signals arrive at R_1 from the direct line-of-sight path and a path due to reflector specified by the tuple $(\theta_{t1,i}, \phi_{t1,i})$, where $\theta_{t1,i}$ denotes the angle of departure of the signal from interferer I_t and $\phi_{t1,i}$ denotes the angle of arrival of the signal at receive antenna R_1

Using sine rule on triangle PI_tQ :

$$\begin{aligned} c_{t1,i} &= PI_t = \frac{\sin\left(\frac{\phi_{t1,i} + \theta_{t1,i}}{2}\right)}{\sin\left(\frac{\phi_{t1,i} - \theta_{t1,i}}{2}\right)} = \frac{d_{t1,0} \sin(\phi_{t1,i}) \sin\left(\frac{\phi_{t1,i} + \theta_{t1,i}}{2}\right)}{\sin(\phi_{t1,i} + \theta_{t1,i}) \sin\left(\frac{\phi_{t1,i} - \theta_{t1,i}}{2}\right)} \\ &= \frac{d_{t1,0} \sin(\phi_{t1,i})}{2 \cos\left(\frac{\phi_{t1,i} + \theta_{t1,i}}{2}\right) \sin\left(\frac{\phi_{t1,i} - \theta_{t1,i}}{2}\right)} = \frac{d_{t1,0} \sin(\phi_{t1,i})}{\sin(\phi_{t1,i}) - \sin(\theta_{t1,i})} \end{aligned}$$

Now suppose the receive antenna R_1 is initially at the origin $(0,0)$. Then the transmitter is at a coordinate $(-d_{t1,0}, 0)$. Let us denote I'_t as the symmetric point of the transmitter about the reflector. Note that since $PI_t = PI'_t$, $PI_t + PR_1 = PI'_t + PR_1 = I'_tR_1$. Hence the length of the reflected path to any receive antenna coordinate R_1 is I'_tR_1 .

Clearly $\angle I'_tI_tR_1 = \pi/2 - \frac{(\phi_{t1,i} - \theta_{t1,i})}{2} = \pi/2 - \beta_{t1,i}$. From $\triangle QI_tS$ we have $SI_t = SI'_t = c \sin \beta_{t1,i}$. Hence $I'_tI_t = 2c_{t1,i} \sin \beta_{t1,i}$. Hence the coordinate of I'_t is $(2c_{t1,i} \sin^2 \beta_{t1,i} - d_{t1,0}, 2c_{t1,i} \sin \beta_{t1,i} \cos \beta_{t1,i}) = (2c_{t1,i} \sin^2(\beta_{t1,i}) - d_{t1,0}, c \sin(2\beta_{t1,i}))$.

Thus the length of the reflected path when the receive antenna is moved to any position $(\delta x, \delta y)$ is:

$$d'_{t1,i} = \sqrt{(2c_{t1,i} \sin^2(\beta_{t1,i}) - d_{t1,0} - \delta x)^2 + (c_{t1,i} \sin(2\beta_{t1,i}) - \delta y)^2}$$

where, $c_{t1,i} = d_{t1,0} \sin \phi_{t1,i} / (\sin \phi_{t1,i} - \sin \theta_{t1,i})$ and $\beta_{t1,i} = (\phi_{t1,i} - \theta_{t1,i})/2$.

B. MINIMIZING INTERFERENCE TO NOISE

Consider two transmitters and a single two-antenna receiver. Let h_{ij} denote the channel from transmitter i to receive antenna j . Our goal is to minimize the interference to noise ratio (INR) for alignment, given by:

$$INR = \frac{1}{\sigma^2} \min_{\alpha, \beta} |h_{11}\alpha + h_{12}\beta|^2 + |h_{21}\alpha + h_{22}\beta|^2$$

where $|\alpha|^2 + |\beta|^2 = 1$ and σ^2 is the noise power.

Without loss of generality, we can choose α to be a positive real, and $\beta = \sqrt{1 - \alpha^2} e^{j\theta}$, where $\theta \in [-\pi, \pi]$. Hence, our goal is to minimize:

$$\begin{aligned} &|h_{11}\alpha + h_{12}\beta|^2 + |h_{21}\alpha + h_{22}\beta|^2 \\ &= (|h_{11}|^2 + |h_{21}|^2)\alpha^2 + (|h_{12}|^2 + |h_{22}|^2)|\beta|^2 \\ &\quad + 2\text{Re}\{\alpha\beta^*(h_{11}h_{12}^* + h_{21}h_{22}^*)\} \\ &= (|h_{12}|^2 + |h_{22}|^2) + \alpha^2(|h_{11}|^2 - |h_{12}|^2 + |h_{21}| - |h_{22}|^2) \\ &\quad + 2\alpha\sqrt{1 - \alpha^2}\text{Re}\{(h_{11}h_{12}^* + h_{21}h_{22}^*)e^{-j\theta}\} \end{aligned}$$

We choose $\theta = \pi + \angle(h_{11}h_{12}^* + h_{21}h_{22}^*)$, so as to minimize the third term. Hence, we need to minimize:

$$\begin{aligned} &(|h_{12}|^2 + |h_{22}|^2) + \alpha^2(|h_{11}|^2 - |h_{12}|^2 + |h_{21}| - |h_{22}|^2) \\ &\quad - 2\alpha\sqrt{1 - \alpha^2}|h_{11}h_{12}^* + h_{21}h_{22}^*| \end{aligned}$$

Let $A = (|h_{11}|^2 - |h_{12}|^2 + |h_{21}|^2 - |h_{22}|^2)$, $B = |h_{11}h_{12}^* + h_{21}h_{22}^*|$ and $C = (|h_{12}|^2 + |h_{22}|^2)$. Hence, we wish to minimize:

$$INR = \frac{1}{\sigma^2} (A\alpha^2 - 2B\alpha\sqrt{1 - \alpha^2} + C) \quad (11)$$

To do so, we set its derivative with respect to α to zero, i.e.

$$\begin{aligned} 2A\alpha - 2B\sqrt{1 - \alpha^2} + 2B\alpha^2/\sqrt{1 - \alpha^2} &= 0 \\ 2A\alpha\sqrt{1 - \alpha^2} - 2B(1 - 2\alpha^2) &= 0 \\ A^2\alpha^2(1 - \alpha^2) &= B^2(1 - 2\alpha^2)^2 \\ \alpha^4(A^2 + 4B^2) - \alpha^2(A^2 + 4B^2) + B^2 &= 0 \end{aligned}$$

The above equation is quadratic in α^2 . Solving, we have:

$$\alpha = \sqrt{\frac{1}{2} - \frac{A}{2\sqrt{A^2 + 4B^2}}}$$

Substituting α in Eqn. 11:

$$\begin{aligned} INR &= \frac{1}{\sigma^2} \left(A \left(\frac{1}{2} - \frac{A}{2\sqrt{A^2 + 4B^2}} \right) - 2B\sqrt{\frac{1}{4} - \frac{A^2}{4(A^2 + 4B^2)}} + C \right) \\ INR &= \frac{1}{\sigma^2} \left((A/2 + C) - \frac{A^2}{2\sqrt{A^2 + 4B^2}} - 2B^2\sqrt{\frac{1}{A^2 + 4B^2}} \right) \\ INR &= \frac{1}{\sigma^2} \left((A/2 + C) - \frac{1}{2}\sqrt{A^2 + 4B^2} \right) \end{aligned}$$

Substituting for A , B and C in the above equation, we have:

$$\begin{aligned} INR &= \frac{1}{2\sigma^2} \left(|h_{11}|^2 + |h_{21}|^2 + |h_{12}|^2 + |h_{22}|^2 \right) - \\ &\quad \frac{1}{\sigma^2} \sqrt{\left| \frac{|h_{11}|^2 + |h_{21}|^2 - |h_{12}|^2 - |h_{22}|^2}{2} \right|^2 + |h_{11}h_{12}^* + h_{21}h_{22}^*|^2} \end{aligned}$$

Which precisely corresponds to Eqn. 7.