# Cross-Cell DoF Distribution: Combating Channel Hardening Effect in Multi-Cell MU-MIMO Networks

Xiufeng Xie and Xinyu Zhang
University of Wisconsin-Madison
{xiufeng, xyzhang} @ece.wisc.edu

Eugene Chai
NEC Labs America
eugene@nec-labs.com

## ABSTRACT

Equipped with the multi-user MIMO (MU-MIMO) technology, a WiFi access point (AP) with $M$ antennas can achieve $M$ degrees-of-freedom (DoF) in theory. Existing MU-MIMO protocols strive to maximize DoF usage by serving $M$ users simultaneously. In this paper, through a MU-MIMO testbed measurement, we found that the correlation between users' channels can severely compromise network throughput under a DoF-maximizing strategy. To combat this problem, we propose Kardia that judiciously distributes the DoF to least-correlated users within an AP's cell, and coordinates neighboring APs such that each can serve the best set of users while nulling mutual interference. The key challenge in Kardia lies in efficiently determining which set of users to serve without knowing the channel state from all of them. We propose a lightweight mechanism that enables APs to collaboratively infer the correlation between users through precoded probing packets. Our analysis shows that Kardia's DoF distribution framework has a provable performance guarantee. Our testbed implementation further reveals that Kardia can significantly boost the performance of multi-cell MU-MIMO networks, in contrast to legacy protocols like 802.11ac.

## Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless Communications*

## General Terms

Algorithms, Design, Theory, Performance

## Keywords

Multi-user MIMO; MU-MIMO; 802.11ac; Channel Hardening

## 1. INTRODUCTION

Multi-user MIMO (MU-MIMO) is a hallmark of the recent WLAN standard 802.11ac, with which an $M_{tx}$-antenna AP can simultaneously transmit up to $M_{tx}$ data streams to different users, thus achieving a maximum *degree-of-freedom* (DoF) of $M_{tx}$. However, exploiting MU-MIMO's full potential is a non-trivial issue in practice. The $M_{tx}$-fold capacity gain is attainable only if channels of all $M_{tx}$ users are *orthogonal* [1] to each other, *i.e.*, the pairwise angles between their *channel vectors*[1] all equal 90°. Consider the example in Figure 1(a), to suppress the mutual interference, the AP precodes U1's data along a direction $\boldsymbol{w}_1$ orthogonal to U2's channel $\boldsymbol{h}_2$. Consequently, U1 receives the signal $\boldsymbol{h}_1\boldsymbol{w}_1$, *i.e.*, a projection of $\boldsymbol{w}_1$ along $\boldsymbol{h}_1$ as shown in Figure 1(b), and thus the received signal strength is $|\boldsymbol{h}_1||\boldsymbol{w}_1|\sin\theta$. Ideally, when U1 and U2 are fully orthogonal, U1's received signal strength is $|\boldsymbol{h}_1||\boldsymbol{w}_1|$. However, the actual signal strength can be much lower — measurement of real MIMO networks revealed that users' channels are often *correlated*, with pairwise angle $\theta$ concentrated between 20° and 70° [2].

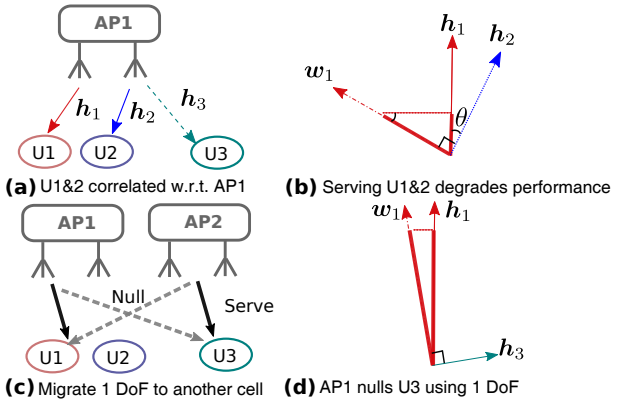


**Figure 1: Impact of user orthogonality in MU-MIMO (a)(b), and how Kardia addresses this problem (c)(d).**

**Channel hardening bounds throughput.** For a given AP, as we add more concurrent downlink users, we increase the probability of serving a pair of users with a small channel angle $\theta$. Thus, MU-MIMO throughput does not scale linearly as more users are added. This so-called *channel hardening* effect has been predicted in theoretical models [3] and also observed in our measurement of indoor MU-MIMO networks (§2). If an AP simply leverages all its spatial DoFs, then the network capacity may be compromised by those strongly correlated users. Hence, the optimal number of DoF used for transmission (*i.e.*, the *usable* DoF) is not always equal to the maximum DoF number supported by the AP. To maximize network throughput, an AP must select a subset of users with optimal mutual channel orthogonality. However, finding such an optimal user set in a large network is an intractable combinatorial problem even for an oracle that knows all users' channels [4].

**Kardia.** We propose a MAC/PHY cross-layer framework, called Kardia[2] to combat the aforementioned channel hardening effect in MU-MIMO networks. The basic principle of Kardia is *cross-cell DoF distribution*. Specifically, we allow an AP to only serve $S \leq M_{tx}$ users that are "most orthogonal" to each other. Its remaining $(M_{tx} - S)$ DoFs are exploited to enable neighboring cells to transmit in parallel. While serving its own users, each AP nullifies its interference to users in other cells (Figure 1(c)(d)). Such a DoF distribution mechanism begs two new questions: *(i) how do we determine the number of usable DoFs, and (ii) how do we distribute these usable DoFs within and across WLAN cells?*

[1]The channel of a single-antenna user can be described as an $M_{tx}$-element *channel vector*, comprised of its channel gain *w.r.t.* each of the AP's transmit antennas.

[2]"Kardia" is a Greek word meaning "Heart", the core organ that distributes blood flows.


*MobiHoc'15,* June 22–25, 2015, Hangzhou, China.

http://dx.doi.org/10.1145/2746285.2746290.

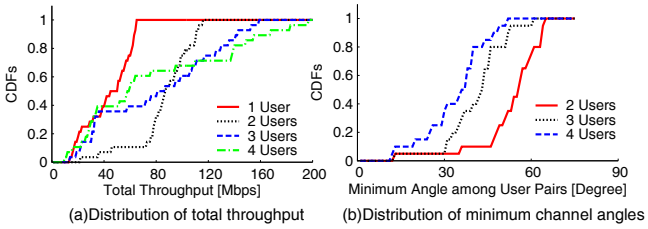


**Figure 2: Network throughput and channel angle when a 4-antenna AP serving from 1 to 4 single-antenna users.**

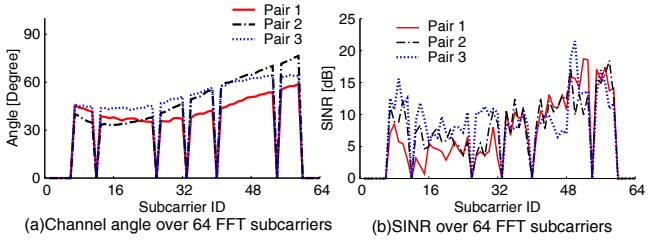


**Figure 3: SINR and channel angle over 64 subcarriers. 0-notches correspond to omitted pilot subcarriers.**

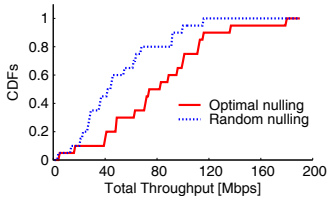


**Figure 4: Randomly choosing users for concurrent transmission does not lead to optimal throughput.**

Ideally, given the channel vectors or channel state information (CSI) of *all users*, an AP can search for the user combination with maximum downlink capacity [4–6]. However, in practice, CSI needs to be obtained from users' feedback. Even if the feedback overhead can be reduced based on channel reciprocity [7] or CSI compression [8], the overhead to obtain global CSI still grows linearly with the total number of users, which can easily overwhelm the actual data transmission time [8]. In addition, selecting the APs that yield the highest cooperation gain is equally challenging. The APs' downlink capacity becomes correlated, and searching through all AP combinations in a large multi-cell network is computationally intractable, even for an oracle that knows global CSI. To meet the above challenges, Kardia employs cross-cell DoF distribution, with three key features:

*(a) Iterative DoF distribution.* Kardia employs a lightweight and scalable MAC protocol, which iteratively groups users and APs before each downlink transmission. The protocol ensures that, within each cell, selected users are most orthogonal to each other, and between cells, every selected AP can effectively nullify its own interference to other cells. This protocol proved to achieve at least $50\%$ of the network capacity of an oracle scheduler.

*(b) Multi-cell orthogonality probing (MOP).* Kardia employs a lightweight MOP mechanism to provide channel orthogonality information and facilitate user/AP selection. It is accomplished without extensive user feedback or global CSI. The AP *learns users' orthogonality with partial CSI knowledge no more than that required for MU-MIMO precoding*, thus bounding feedback overhead.

*(c) Orthogonality-aware subcarrier selection (OAS)* allows an AP to turn frequency selectivity into an advantage. It concentrates transmit power on subcarriers that tend to have strong CSI orthogonality when used by selected users.

We prototype Kardia based on the WARP platform. Our testbed first realizes a MU-MIMO OFDM physical layer following the 802.11ac specification. Then, we build the Kardia components on top of the PHY module. Our testbed experiments demonstrate that:

*(i)* Kardia's DoF distribution mechanism effectively picks capacity maximizing users/APs combinations, and can achieve near-optimal performance in practical network topologies.

*(ii)* Kardia has a low overhead. It only collects CSI from selected users rather than all users in the network. As a result, Kardia bears comparable signaling overhead as vanilla 802.11ac, and its overhead is bounded even for a large network with up to 100 users.

*(iii)* In large-scale testbed experiments, Kardia achieves around $2.07\times$ and $1.56\times$ throughput gain over vanilla 802.11ac and state-of-the-art MU-MIMO user selection scheme [9], respectively.

## 2. MOTIVATION

In this section, we motivate Kardia's design components through micro-benchmark measurements on a MU-MIMO testbed.

**Using All DoFs is Ineffective.** Real-world MU-MIMO transmissions are subject to the aforementioned *channel hardening* effect. As a result, the sum-rate of a transmitter does not increase linearly with the number of downstream users. We study the prevalence of this channel hardening effect through a MU-MIMO testbed with a 4-antenna AP and single-antenna users (§6 details our testbed). We conduct the experiments over 20 different topologies. In each topology, we record the total throughput when the AP serves from 1 to 4 users, respectively. Figure 2(a) shows the resulting throughput distribution for different user numbers. We observe that the median throughput doubles from 40 to 80Mbps when we increase the user number from 1 to 2. However, it *decreases* as we further increase the user number to 4 — *For almost half of the cases, serving 2 users renders a higher sum-rate than serving 4 users.*

Such channel hardening effect roots in the spatial correlation between user's channels, which bottlenecks MU-MIMO capacity [3]. Under the above setting, Figure 2(b) shows the distribution of minimum pairwise angles between users (*i.e.*, the indicator for spatial channel correlation, as discussed in §1). Note that if we blindly select $4$ users for a MU-MIMO transmission, the minimum pairwise angle is below $35^\circ$ for $50\%$ of the cases, thus crippling the achievable throughput and explaining the results in Figure 2(a).

**Using All Subcarriers is Ineffective.** In a frequency-selective channel, the CSI varies across subcarriers. Hence, the degree of channel orthogonality between a pair of users will also vary across subcarriers. Figure 3(a) shows the angle between three distinct user pairs over $64$ subcarriers in a 20MHz band. Observe that for a given pair of users, the channel angle between them can vary significantly across subcarriers. Consequently, the SINR and associated throughput over the subcarriers can vary substantially, as seen in Figure 3(b). Since the sum power of subcarriers is always bounded by the transmitter's total power budget, it will be beneficial to transmit only over subcarriers with high user orthogonality.

**Blind Interference Nulling is Ineffective.** When an AP serves its own users while nulling its interference to other-cells' active users, interestingly, the AP's choice of other-cell-users to be nulled also affects the throughput of its own users. Figure 4 compares the total throughput achieved with optimal and random nulling. In each scenario, a 4-antenna AP selects either 2 or 3 MU-MIMO users and uses the remaining DoF for nulling additional users. For optimal nulling, we select the additional users that are most orthogonal to the users to be served, while for random nulling, nulled user(s) are selected at random. We observe that optimal nulling almost doubles the median throughput among the MU-MIMO users. Hence, we must carefully choose the optimal nulling candidates when performing DoF distribution in Kardia.

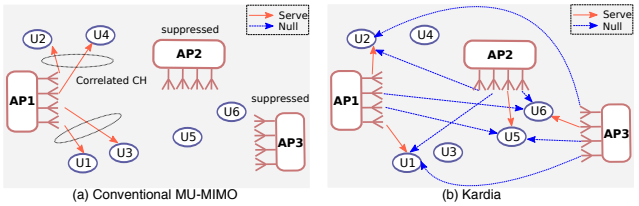


**Figure 5: Final schedule in an example topology.**

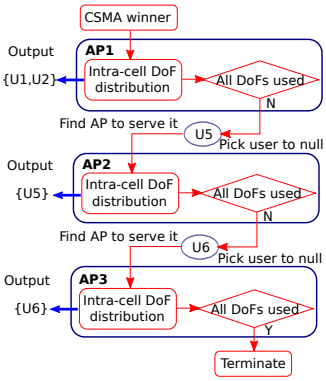


**Figure 6: Kardia operation flow for the example in Fig. 5.**

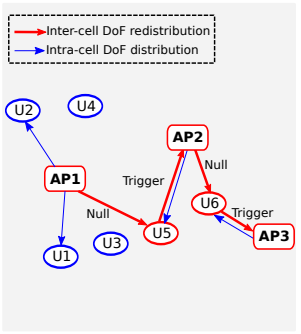


**Figure 7: Iterative search path in user/AP selection.**

# 3. Kardia: A CASE STUDY

To understand how Kardia works, let us first focus on an example network shown in Figure 5, where each AP has $M_{tx} = 4$ antennas and serves *single-antenna* users. We assume a single contention domain, *i.e.*, all users can be reached by all APs in this network. We further assume that the channels of U1 and U3 *w.r.t.* AP1 are correlated (*i.e.*, the channel angle between AP1→U1 and AP1→U3 is close to $0^\circ$). Similarly, U2 and U4 are correlated as well. U5 is uncorrelated with all other users, but has the lowest channel gain since it is furthest away from AP1.

Suppose AP1 wins channel contention, vanilla 802.11ac will allow AP1 to exhaust the DoFs and serve *all* 4 downlink users, causing low performance. In contrast, Kardia allows AP1 to identify the usable DoFs and serve the best un-correlated *subset* of users (say {U1,U2} in this example). AP1 then uses its remaining 2 DoFs to nullify its interference to users that may be served by other APs, thus enabling concurrent transmissions of multiple APs (Fig. 5(b)).

To maximize the capacity of concurrent downlink transmission, Kardia optimizes the AP/user selection as follows (Fig. 6 and 7).

**Intra-cell DoF distribution.** Kardia is initiated by AP1 immediately after it gains channel access. AP1 first instantiates *intra-cell DoF distribution* (§4.1.2) to *iteratively* add users into its MU-MIMO transmission group to maximize intra-cell downlink capacity. The first user (U1 in this example) to add is randomly picked with a probability ensuring proportional fairness (§4.5). Then U2 is added since its channel is the most orthogonal to U1's.

At this stage, combining any more user along with U1 and U2 will result in an increased channel condition number [1] and thus degrades downlink capacity. AP1 thus stops adding downlink users with 2 unallocated DoFs "blocked" by the channel hardening effect.

**Inter-cell DoF distribution.** To leverage the unallocated DoFs, after AP1 terminates its intra-cell DoF distribution, Kardia selects a *bridge client* for AP1 to nullify. It employs multi-cell orthogonality probing (MOP, §4.2) to ensure that the bridge client can be simultaneously served by a neighboring AP, and that the resulting inter-cell interference can be nulled by AP1.

In this example, AP1 selects U5 as the *bridge client*. Although U5 is uncorrelated to U1 and U2, AP1 does not directly serve them together as this will degrade downlink capacity[3], however, steering the downlink beams away from U5 will only cause trivial signal power loss on U1 and U2 since they are uncorrelated, hence AP1 can nullify its interference to U5 without affecting the performance of U1 and U2. U5 then triggers a new AP (AP2 in this case) to serve it as shown in Figure 7.

The newly activated AP2 then continues to run its own intra-cell DoF distribution. Note that each AP needs to nullify its interference to other APs' selected users. Since U3 and U4 are correlated to AP1's selected users (U1 and U2), they may reduce total network capacity when selected by AP2. Therefore, AP2 excludes U3 and U4, and instead selects a bridge client (U6) uncorrelated to all currently selected users *w.r.t.* all current APs. The bridge client U6 then triggers a new AP (AP3) to serve it.

***Termination of Kardia.*** Such iterative user-AP selection process continues until one of the APs has no DoF left. In the final schedule shown in Figure 5(b), this happens after U6 is paired with AP3 and consequently each AP has 4 users to serve/null in total. Thereafter, each AP activates its MU-MIMO PHY-layer transmission module to serve its own users while nulling other APs' users.

It is worth noting that although we introduce Kardia's work flow here based on single-antenna users, our design can be easily extended to multi-antenna users. In effect, we can treat each antenna on a multi-antenna user as a separated "single-antenna user".

**Algorithm 1** Kardia's outline.

1. The *lead AP* wins CSMA channel contention
2. Add the *lead AP* in the set of selected APs $\mathcal{A}$
3. The *lead AP* picks its first user and gets CSI feedback
4. **While** the minimum Tx DoF among all APs > 0
5. /*Set current scheduling AP $j$ to be the newly added AP*/
6. AP $j$ performs *intra-cell DoF distribution* (Algorithm 2)
7. AP $j$ determines its set of users $\mathcal{U}_j$ to serve
8. Any other selected AP $k$ updates its set of users to null $\mathcal{N}_k$
9. **If** remaining DoF $D_j > 0$
10. AP $j$ performs *inter-cell DoF distribution* (Algorithm 3)
11. Add the newly selected AP $j'$ to the set $\mathcal{A}$
12. **Output**: the set of all selected APs $\mathcal{A}$
    the users to be served by each AP $\{\mathcal{U}_1, \mathcal{U}_2, ..., \mathcal{U}_{|\mathcal{A}|}\}$
    the users to be nulled by each AP $\{\mathcal{N}_1, \mathcal{N}_2, ..., \mathcal{N}_{|\mathcal{A}|}\}$

# 4. Kardia DESIGN

## 4.1 Iterative DoF Distribution

### 4.1.1 An outline of operations

We now generalize the above example and detail Kardia's iterative DoF distribution protocol. The protocol outputs a near-optimal set of APs/users that maximizes the sum capacity of concurrent downlink transmissions, while minimizing costly CSI feedback. Algorithm 1 summarizes the major steps. Here we assume homogeneous APs, each with $M_{tx}$ transmit antennas[4].

The DoF distribution protocol is activated whenever an AP gains channel access following CSMA. This AP is referred to as the *lead AP*. The lead AP then announces its intention for MU-MIMO trans-

[3]In MU-MIMO, the user with a much worse channel gain than others will take most of the transmit power after precoding, nulling this user avoids the problem since no power is allocated to that user.

[4]It is straightforward to extend our algorithm to heterogeneous AP cases, where the total number of DoFs is bounded by the antenna number on the AP that wins the CSMA contention first. This is because Kardia basically transfers the DoFs of this first AP to other cells when the DoFs are "blocked" by the channel hardening effect.

**Algorithm 2** Intra-cell DoF distribution.

1. Input: current AP $j$, CSI of its first user $\boldsymbol{H}_0 = \boldsymbol{h}_{1j}$, its remaining DoF $D_j = M_{tx} - 1$
2. **While** $D_j > 0$
3. AP $j$ precodes and sends the probing frame
4. **If** Any unselected user has orthogonality metric $> \tau_o$
5. Unselected users perform distributed ranking (§4.2).
6. The winner user $k$ feeds back CSI to AP $j$
7. AP $j$ adds user $k$ to the selected group $\mathcal{U}_j$
8. AP $j$ reduces the remaining DoF $D_j = D_j - 1$
9. **Else** AP $j$ times out if no reply. Break.
10. Output: selected user group $\mathcal{U}_j$, remaining DoF $D_j$

mission through an 802.11ac signaling packet called *NDP*. The first user to be served is picked based on proportional fairness (§4.5) and also announced in NDP. The first user $U_i$ then feeds back its CSI, *i.e.*, its channel vector $\boldsymbol{h}_{ij}$ *w.r.t.* the *lead AP* $j$. Thereafter, AP $j$ instantiates *intra-cell DoF distribution* (§4.1.2) to iteratively add users to serve, which ensures that each newly selected user has the highest potential to boost capacity. If an AP terminates its *intra-cell DoF distribution* early and leaves surplus DoFs, Kardia proceeds with *inter-cell DoF distribution* (§4.1.3), allowing the AP to use its surplus DoFs to null additional users so that they can be served by other APs. The intra- and inter-cell DoF distribution iterates until all available APs terminate or all $M_{tx}$ DoFs are distributed.

### 4.1.2 Intra-cell DoF distribution

Algorithm 2 summarizes the *intra-cell DoF distribution* in Kardia, which is designed to achieve the following property:

**Property 1.** *Whenever a new user is added, among all unselected users, it has maximum orthogonality to any existing selected users w.r.t. any existing selected APs.*

Below we introduce the intra-cell DoF distribution design.

**Multi-cell Orthogonality Probing (MOP).** At the start of *intra-cell DoF distribution*, AP $j$ already has its first user to serve. It instantiates Kardia's core component, *i.e.*, multi-cell orthogonality probing, to iteratively add more users. To single out the user mostly orthogonal to currently selected ones, a naive approach is to collect CSI feedback from all unselected users. Yet this causes formidable overhead. Instead, the Kardia AP beamforms a customized *probing frame* toward directions orthogonal to the CSI of any currently selected user. The probing frame enables unselected users to estimate their channel quality and orthogonality to currently selected ones *w.r.t.* AP $j$ based on the received SNR. After hearing the *probing frame*, the unselected users initiate a *distributed ranking*, where the one with the highest orthogonality metric wins and feeds back its CSI to AP $j$. The AP keeps adding users successively by repeating the above steps. We defer the detailed design of probing frame, orthogonality estimation, and distributed ranking to Section 4.2.

**Termination when adding more users is harmful.** To guarantee the channel orthogonality within the selected user set $\mathcal{U}_j$, only users with orthogonality metric higher than a threshold $\tau_o$ are allowed to participate in the distributed ranking after AP $j$'s probing frame. If in a certain iteration, the AP hears no CSI feedback, then all unselected users must have orthogonality metric $< \tau_o$. In this case, AP will time out and terminate intra-cell DoF distribution early, even if it has unallocated DoFs. Otherwise the algorithm iterates until the AP's DoF is exhausted.

**Algorithm 3** Inter-cell DoF distribution.

1. Input: current AP $j$, set of its downlink users to serve $\mathcal{U}_j$
2. AP $j$ broadcasts a notification to pick nulling target
3. Unselected users perform distributed ranking
4. **If** Any unselected user has orthogonality metric $> \tau_n$
5. User $l$ most orthogonal to $\mathcal{U}_j$ wins and feeds back its CSI
6. Every selected AP $\in \mathcal{A}$ adds user $l$ in its nulling set $N$
7. CSI feedback from user $l$ triggers additional AP selection
8. Unselected AP evaluates orthogonality if it joins $\mathcal{A}$
9. AP $j'$ who ensures the best user orthogonality joins $\mathcal{A}$
10. Output: $U_l$ to be nulled by AP $j$, the new AP $j'$ to serve $U_l$

### 4.1.3 Inter-cell DoF distribution

Algorithm 3 summarizes the *inter-cell DoF distribution* in Kardia, which is geared toward the following goal:

**Property 2.** *Whenever a new AP is added, among all unselected APs, it maximizes the channel orthogonality of all existing selected users (which it needs to null).*

Below we detail the inter-cell DoF distribution design.

**Step1: Selecting a user to null.** Suppose an AP $j$ terminates intra-cell DoF distribution early and leaves surplus DoFs due to channel hardening, it then starts the *inter-cell DoF distribution* by picking a user to null, *i.e.*, the *bridge user*. The choice here is as important as selecting a user to serve, since nulling an improper user may ruin the performance of existing selected users (§2).

To pick the optimal user to null, AP $j$ broadcasts a short frame to notify other users/APs that it has already terminated intra-cell DoF distribution with DoFs left. All unselected users overhearing this notification then conduct a special distributed ranking based on an orthogonality metric estimated from AP $j$'s last probe (detailed in §4.2). Specifically, to ensure nulling the bridge user will not significantly affect the downlink capacity of AP $j$, among the unselected users, only those with orthogonality metric greater than a threshold $\tau_n$ can join this distributed ranking. The user $U_l$ most orthogonal to AP $j$'s existing users wins and feeds back its CSI.

**Step2: Selecting new AP.** The bridge user $U_l$ then elects a new AP $j'$ to serve it. This AP should have the highest potential capacity when serving the users nulled by previous AP(s) and nulling users served by previous AP(s). To gauge its potential, each AP who overheard the CSI feedback from $U_l$ evaluates the *condition number* of the channel matrix containing all CSI feedbacks it has overheard up to now. A smaller condition number implies a smaller correlation between channels and thus higher potential capacity [1]. These candidate APs report their evaluated condition number to a central coordinator via a wireline control channel as in [10]. The AP with the smallest condition number is then selected.

**Step3: User selection for the new AP.** The newly added AP then continues to run its own intra-cell DoF distribution procedure. Note that all previous AP(s) need to nullify interference to the new AP's selected users. Essentially, each AP schedules a set of users to serve/null, and each user can join multiple AP's scheduled sets (served in one set, nulled in all other sets). The challenge here is to let each user ensure the channel orthogonality in every scheduled set it takes part in. Kardia adopts a simple but effective solution: if a user has found itself correlated to any selected user in an AP's intra-cell DoF distribution iteration, it will simply ignore future probes from new APs and hence will not be selected by any of them for this downlink MU-MIMO transmission.

**Step4: Termination of DoF distribution.** Such joint user-AP selection iterates until the network's DoF is exhausted, *i.e.*, the total number of users to serve reaches $M_{tx}$. The final output is a downlink transmission schedule including a set of APs, each to serve its own set of users and null those belonging to other APs. Afterwards, each AP further examines the user coupling across its bandwidth via an *orthogonality-aware subcarrier selection* (OAS) scheme (§4.3). Finally, each AP runs Zero-Forcing Beamforming (ZFBF), a standard MU-MIMO precoding algorithm [1], to serve

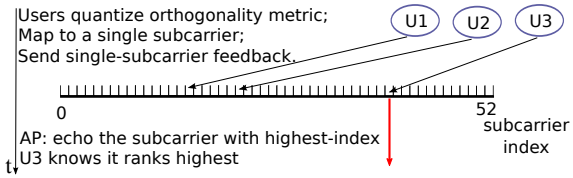


**Figure 8: Distributed ranking in MOP.**

its users while nulling other APs' users. Note that the CSI used in ZFBF is exactly that collected during Kardia's AP/user selection. No extra CSI feedback is required. Therefore, we have:

**Property 3.** *Kardia does not require global CSI knowledge. The number of users who feed back their CSI equals the number of users served in the subsequent MU-MIMO transmission, which is typically much smaller than the total user population in the network.*

## 4.2 Multi-cell Orthogonality Probing (MOP)

We now detail the basic primitives in multi-cell orthogonality probing (MOP), the key PHY-layer operation to enable the DoF distribution protocol discussed above.

**Probing frame.** Before adding a user to serve/null, the AP sends a probing frame to help unselected users assess their orthogonality with already selected ones. The basic idea is to probe channel directions that are orthogonal (in the channel vector space) to channels of already selected users. Each channel direction carries a training sequence embedded inside the frame's preamble. The probing frame design follows existing single-cell MU-MIMO signaling schemes [9]. We thus omit more details here.

**Orthogonality metric.** A user's orthogonality metric in MOP reflects its potential contribution to downlink capacity when grouped with selected users, and served/nulled by existing APs. More specifically, it should reflect the user's: *(i)* channel quality *w.r.t.* the probing AP, and *(ii)* channel orthogonality to already selected users *w.r.t.* the probing AP. Since the probing directions are orthogonal to channels of selected users, the *maximum SNR among all probing directions* exactly meets the two requirements above. A user can evaluate its SNR values along each probing direction by decoding the corresponding known training sequence. The training sequence with highest SNR corresponds to the most orthogonal direction. We screen users with orthogonality metric lower than $\tau_o$ (§4.1.2), a threshold set to 6 dB (SNR threshold for the lowest modulation level). We validate this orthogonality metric design in §7.1.

Note that for the *bridge client*, MOP should only gauge its channel *orthogonality* to already selected users, and isolate the effect of channel quality (reflected by link SNR). We thus subtract the bridge client's link SNR from the orthogonality metric (only when selecting the bridge client). The link SNR is estimated directly from the probing AP's preamble following [8]. Correspondingly, the threshold $\tau_n$ (§4.1.3) equals $\tau_o$ subtracting the link SNR.

**Distributed Ranking.** MOP's distributed ranking mechanism is inspired by recent work in frequency-domain contention in CSMA [11, 12]. Instead of random access, however, it is used to single out the user with highest orthogonality metric. Specifically, after receiving a probing frame, each unselected user $k$ computes and quantizes its orthogonality metric $g_k$ into integers within $[0, G_m]$. The quantized value $\hat{g}_k$ is then mapped to index of OFDM subcarriers ranging from 0 to $G_m$. Then, the user sends an OFDM symbol with only the $\hat{g}_k$-th subcarrier carrying a non-zero data symbol $(1+j)$, as illustrated in Figure. 8. After receiving the feedback, the AP employs energy detection over frequency domain to identify the active subcarrier with highest index. It considers a subcarrier to be active if the energy level is $\tau$ dB above the noise floor (estimated by the average energy of unpopulated guardband subcarriers). An empirical value of $\tau = 3$ dB is found to achieve high detection

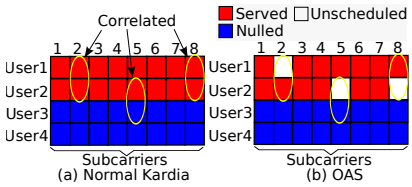


**Figure 9: Simple example for OAS.**

accuracy in practice. Afterwards, the AP echoes back an OFDM symbol by activating the highest-index subcarrier, so that the corresponding user knows it is selected.

A notable feature of our distributed ranking design is its constant overhead irrespective of the user population. But when the users have similar quantized orthogonality metric, a collision may occur. MOP employs several counter-measures: *(i)* Only users with orthogonality metric above a lower limit $\tau_o$ can participate in the distributed ranking, hence the spectrum becomes less congested. *(ii)* Kardia also employs an upper limit corresponding to the maximum SINR that AP observes among all users it can serve. These two limits bound the range of orthogonality metric to quantize, resulting in a high resolution after quantization. We choose $[7, 30]$ dB as the default range. With 52 subcarriers (excluding guardband) in 802.11, the resolution becomes 23/52 = 0.44dB. In practice, this resolution ensures that users in a medium sized network can seldom share the same highest quantized orthogonality metric (§7.1). A larger number of subcarriers can lead to even higher resolution. *(iii)* In the rare event when collision occurs, the AP may receive CSI feedbacks from multiple users and cannot resolve them. It will terminate the probing and instruct all already selected users/APs to start MU-MIMO transmission, so as not to waste the channel time.

## 4.3 Orthogonality-Aware Subcarrier Selection

To bound the signaling overhead, Kardia selects users based on an orthogonality metric that averages the heterogeneous channels of different subcarriers and thus may not properly capture the frequency diversity. We resolve this issue through Orthogonality-Aware Subcarrier Selection (OAS), a post-schedule refinement to Kardia's output. As illustrated in Figure 9, OAS excludes a user on a certain subcarrier if it bears weak orthogonality to other users there. This brings two benefits: *(i)* It focus more signal power on subcarriers with good channel orthogonality, thus improving SNR. *(ii)* Even if such additional power boost is small, eliminating some selected users on the "bad orthogonality" subcarriers is still critical since they will otherwise induce high bit error rates to other users who could have much better performance without them.

To execute OAS, for each subcarrier, an AP computes the Hermitian angle between each pair of users that it needs to serve/null. Since the users already fed back their CSI, the AP has enough CSI knowledge to perform this operation. If the angle of a pair is less than a threshold $\theta_a$, one of the users will be excluded on this subcarrier. The user with bad channel orthogonality to more users will be excluded with higher priority. We empirically set $\theta_a = \pi/6$, a relative angle where two users only achieve half of the effective transmit power compared with the fully orthogonal case (§1).

## 4.4 SINR estimation for MU-MIMO downlink

Finally, before MU-MIMO transmission, each AP needs to determine the modulation level for each user. The key challenge here lies in the link quality estimation. Each MU-MIMO transmission may involve a different set of user combinations, whose SINR depend their orthogonality, and therefore, history based adaptation no longer works. Ideally, the AP can model the SINR following the well-known MU-MIMO capacity equation [1]. However, it is non-trivial to obtain the noise level in practice. Simply measuring the

"idle" energy level fails to take into account the channel noise (*e.g.*, caused by imperfect channel estimation) which may vary among different APs. It is possible to directly measure and feedback SINR by individual test transmission, but this incurs enormous overhead.

Kardia employs a lightweight method to obtain SINR. In a user $U_i$'s CSI feedback, it piggybacks its orthogonality metric, expressed as a 16-bit float point number in dB scale. The orthogonality metric can be modeled as: $SINR_i^p = S_i^p/(I_i^p + N_i)$, where $S_i^p = \boldsymbol{h}_i\boldsymbol{p}$ is the received power of $U_i$ along probing direction $\boldsymbol{p}$ (the direction where it experiences the highest SINR) and $N_i$ is the noise level. $I_i^p$ is the leakage interference from other directions to direction $p$. Note that leakage interference exists because the probing packet is not directly beamformed to $U_i$. Meanwhile, the AP computes the signal to leakage-interference ratio: $SLR_i^p = S_i^p/I_i^p$, based on the fact that $I_i^p = \sum_{\boldsymbol{q}\neq\boldsymbol{p}} \boldsymbol{h}_i\boldsymbol{q}$.

Since both $SINR_i^p$ and $SLR_i^p$ are ratio between power levels, we can circumvent the unknown scalars introduced by the transceivers' power amplifiers and other analog components. After receiving the feedback, the AP can obtain the noise level: $N_i = (SINR_i^p)^{-1} - (SLR_i^p)^{-1}$. Then, the expected SINR of $U_i$ if served by AP $j$ can be modeled as $SINR_i = \gamma_j\boldsymbol{h}_i\boldsymbol{w}_i/N_i$. Note that ZFBF dictates that $I_i = 0$. $\gamma_j$ is the power scalar that respects the AP's power budget. When OAS is enforced, $\gamma_j$ is scaled by the ratio between total number of subcarriers and actual number that carriers data.

## 4.5 Other Practical Issues

**Fairness control.** Kardia employs a randomized algorithm to arbitrate proportional throughput fairness among users. Specifically, in the beginning of MOP, the lead AP selects the first user with a probability weighted by users' historical throughput record, which the AP learns from ACKs. The probability that user $k$ is selected as the first user is $p_k = (1/\overline{\phi_k})/\sum_{j\in\mathcal{K}}(1/\overline{\phi_j})$, where $\overline{\phi_k}$ is its average throughput and $\mathcal{K}$ is the set of all users. Such a scheduling mechanism has proven to achieve proportional fairness [1].

**Overhead of Kardia.** Owing to its light-weight MOP, Kardia only adds marginal overhead over vanilla 802.11ac. In particular, *its number of probing/feedback rounds is bounded by $M_{tx}$, irrespective of the network size.* In each round, *(i)* The probing frame contains the 802.11ac STF, LTF and VHT-LTF preambles, $20\mu s$ in total. *(ii)* The implicit CSI feedback contains STF and LTF, $16\mu s$ in total). *(iii)* The distributed ranking consists of two OFDM symbols ($8\mu s$) plus a TX/RX switching time ($9\mu s$ in 802.11), costing $17\mu s$ in total. When $M_{tx} = 4$, MOP's overhead is approximately $(20+16)+(20+16+17)\times(M_{tx}-1) = 195\mu s$ (note that the first round does not entail distributed ranking). Compared with vanilla 802.11ac, the outstanding overhead of Kardia lies in the frequency-domain distributed ranking, which is only $17(M_{tx}-1)\mu s$. In practice, frame aggregation is commonly enforced on MU-MIMO transmissions to amortize the overhead, which is evaluated in §7.1.

**Multiple contention domains.** Our previous description of Kardia has been restricted to a single contention domain where APs/-users can sense each other. In a large network, the current Kardia design arbitrarily partition it into multiple clusters, each being one contention domain. Only APs/users inside each cluster coordinate and run Kardia. Those in different clusters either follow the contend-and-reserve protocol in 802.11ac, or relies on a TDMA scheduler running by a central coordinator. It is possible to further optimize the clustering as in [13], but this is left as our future work.

**Synchronization.** Kardia's PHY layer leverages two mature communications technologies, including intra-cell MU-MIMO precoding and inter-cell interference nulling. Neither require sophisticated carrier frequency or phase synchronization among distributed APs or clients. The latter does require coarse time synchronization to within an OFDM cyclic prefix length [10], which is 0.8 $\mu s$ in 802.11ac. Existing system prototype has verified the feasibility of interference nulling under such synchronization requirement using commodity WiFi cards [10]. We remark that distributed ranking design in Kardia also requires similar level of coarse synchronization among clients and AP within an OFDM cyclic prefix length.

**Channel coherence time.** It should be noted that multiple rounds of downlink transmissions can follow one DoF distribution, provided that they span much shorter time than the channel coherence time and thus CSI remains stable. Existing experiments [13] observed that for indoor environment with low-mobility and short-distance links (*i.e.*, the targeted scenario of Kardia), MU-MIMO capacity is unaffected as long as CSI is refreshed every 15 ms. In our evaluation, MOP is performed based on a conservative update threshold of 5.5 ms[5], *i.e.*, it is initiated before a MU-MIMO transmission, if more than 5.5 ms has elapsed since the last MOP. This conservative approach tends to amplify the overhead of Kardia.

**Natural regression to 802.11ac mode.** When the APs are sparsely deployed, *i.e.*, we have independent cells without interference to each other, the 802.11ac can simply work in each cell. In this case, Kardia naturally regresses to the vanilla 802.11ac. In the example in Sec. 3, if AP2 does not interfere with the users of AP1, then the CSI feedback of AP1's users will never reach AP2 due to channel reciprocity. Hence AP2 still has its full $M_{tx}$ DoFs to serve its own users. Such graceful regression guarantees Kardia can achieve a throughput no less than vanilla 802.11ac in any scenario.

# 5. Kardia: PERFORMANCE GUARANTEE

Kardia's iterative DoF distribution algorithm tries to maximize user orthogonality in each iteration, without incurring the formidable computational and feedback overhead in an exhaustive search. But how far can it deviate from a global exhaustive search? We now answer this question by lower-bounding Kardia's performance in the basic multi-cell MU-MIMO and interference nulling setup (§4.1).

**(a) DoF distribution as an optimization problem.** We introduce an indicator variable $x_{uj} \in 0, 1$, denoting whether AP $j$ serves user $u$. We assume the network is dense, such that for each user, there are a sufficient number of users (belonging to any APs) that can form an orthogonal group with it. Even under this assumption, it is non-trivial to solve the problem of DoF distribution that maximizes total network capacity, which can be formulated as a non-linear integer program, as follows:

$$\max_{\boldsymbol{x}} \quad g(\boldsymbol{x}) \tag{1}$$

$$\text{s.t.} \quad \textstyle\sum_{j\in\mathcal{A}} x_{uj} \leq 1, \forall u \in \mathcal{U} \tag{2}$$

$$\textstyle\sum_{j\in\mathcal{A}}\sum_{u\in\mathcal{U}} x_{uj} \leq M_{tx} \tag{3}$$

where $\mathcal{A}$ and $\mathcal{U}$ denote set of APs and users to choose from. Eq. (2) says that a user is served by at most one AP, whereas Eq. (3) captaures the DoF constraint, *i.e.*, total users to be served should be no larger than $M_{tx}$, the number of antennas on each AP.

Our objective is to maximize $g(\boldsymbol{x})$, the total network capacity for a given user-grouping and AP selection strategy $\boldsymbol{x}$. Suppose each AP has a power budget of $P$, then under ZFBF precoding, the AP $j$ projects equal power $\gamma_j$ towards each user it serves. Thus, for each AP $j$, we have:

$$\sum_{u\in\mathcal{U}} \frac{\gamma_j x_{uj}}{||\vec{h}_{uj}||^2} = P, \text{and } \gamma_j = \Big(\sum_{u\in\mathcal{U}} ||\vec{h}_{uj}||^{-2} x_{uj}\Big)^{-1} P \tag{4}$$

Then, the total network capacity can be written as:

$$g(\boldsymbol{x}) = \sum_{j\in\mathcal{A}}\sum_{u\in\mathcal{U}} x_{uj}\log(1+\frac{\gamma_j}{N_0}) \approx \sum_{j\in\mathcal{A}}\sum_{u\in\mathcal{U}} x_{uj}\log\frac{\gamma_j}{N_0} \tag{5}$$

[5]5.5 ms is the maximum allowable frame duration in 802.11ac.

**(b) Submodularity of objective function.**

*Definition 1: Submodular function.* If $\Omega$ is a set and $2^{\Omega}$ its power set, then a set function $f$ is submodular *w.r.t.* $\Omega$ if: For every $X, Y \subseteq \Omega$ with $X \subseteq Y$ and every $e \in \Omega \backslash Y$, we have $f(X \cup e) - f(X) \geq f(Y \cup e) - f(Y)$.

Submodularity is a diminishing returns property, stating that adding an element to a smaller set helps more than adding it to a larger set. In our problem, intuitively, as the user set becomes larger, orthogonality is more likely to become worse when a new user is added (§2). Thus, the objective function is likely to be submodular. We formalize this intuition as follows.

**Lemma 1.** *The function $g(\boldsymbol{x})$ is submodular.*

*Proof sketch:* Define $\Omega = \{(u, j) : u \in \mathcal{U}, j \in \mathcal{A}\}$, then $\boldsymbol{x}$ is its power set. For a given set $U^0 = \{(u, j) \in \Omega : x_{uj} = 1\}$, suppose one more element $(u', j') \notin U^0$ is added to $U^0$. Then, for the same element $(u', j')$ added to $V^0$, a subset created by removing one element from $U^0$, we can easily prove that $\left[g\left(U^0 \cup (u', j')\right) - g(U^0)\right] - \left[g\left(V^0 \cup (u', j')\right) - g(V^0)\right]$ is non-negative. By induction, we can prove $g(\boldsymbol{x})$ satisfies the submodularity definition. We ommit the details due to space constraint. □

**(c) Worst-case performance bound.** Problem (1) resembles the classical problem of maximizing a submodular set function, particularly the maximum general assignment problem [14, Ch. 3]. Approximation ratio of assignment problems with exclusive constraint like (2) is well known (see a survey in [14] and recent works, *e.g.*, [15, 16]). However, the DoF constraint (3), which limits the cardinality of set $\boldsymbol{x}$, is unique in our setting. To our knowledge, performance bound of iterative solutions to such a problem has not been analyzed in prior work.

**Theorem 1.** *The performance of Kardia's iterative user/AP selection mechanism is at least $\frac{1}{2}$ of the global optimum.*

*Proof:* Our performance analysis extends the lower-bounding techniques used in optimizing submodular functions defined over a *partition matroid* [17]. We partition the ground set $\Omega$ into $|\mathcal{U}|$ sets, with the $u$-th set being: $\mathcal{S}_u = \{(u, j) : j \in \mathcal{A}\}$, *i.e.*, collection of links with the same user $u$. Then the set of sets $\mathcal{F} = \{\cup_{u=1}^{|\mathcal{U}|} |\mathcal{F}_u$, where $\mathcal{F}_u \subseteq \mathcal{S}_u$ and $|\mathcal{F}_u| \leq 1, \forall u = 1, 2, \cdots, |\mathcal{U}|\}$ becomes a partition matroid. Note that the constraint of $|\mathcal{F}_u| \leq 1$ is equivalent to Eq. (2).

We now show that the cardinality constraint (3) reduces Kardia's iterative algorithm to an early-termination version of the classical greedy solution to optimizing a partition matroid [17], but the approximation ratio of $\frac{1}{2}$ still holds.

Denote $L^*$ as the set of user-AP pairs in the optimal solution to (1), and $L$ as Kardia's algorithm. Let $\rho_e(A)$ be the increment to the objective function when a new element $e$ is added to the solution set, *i.e.*, $\rho_e(A) = g(A \cup e) - g(A)$. Following a well known property of submodular set functions [17], we have:

$$g(L^*) \leq g(L) + \textstyle\sum_{l \in L^* \backslash L} \rho_l(L) \tag{6}$$

Let $e$ be the element in $L$ that has the minimum increment when it is added. $L_e$ denotes the solution set right before $e$ is added. The submodularity of $g(\cdot)$, along with the selection step of Kardia, implies that:

$$\textstyle\sum_{l \in L^* \backslash L} \rho_l(L) \leq \sum_{l \in L^* \backslash L} \rho_l(L_e) \leq \sum_{l \in L^* \backslash L} \rho_e(L_e) \tag{7}$$

$$\text{Further, } \textstyle\sum_{l \in L^* \backslash L} \rho_e(L_e) = |L^* \backslash L| \rho_e(L_e) \tag{8}$$

$$\leq M_{tx} \rho_e(L_e) \leq \textstyle\sum_{l \in L} \rho_l(L_l) \tag{9}$$

$$= g(L) \tag{10}$$

Inequality (9) follows the definition of $e$. Combining Eq. (10) with (6), we have $g(L) \geq \frac{1}{2} g(L^*)$, which completes the proof. □

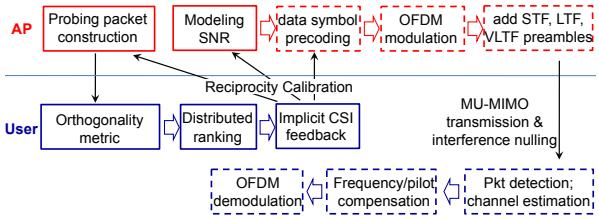


**Figure 10: Implemented modules in Kardia. Boxes with broken lines represent 802.11-compliant PHY modules.**

# 6. IMPLEMENTATION

We build a prototype of Kardia using the WARP software radio platform. Figure 10 shows the basic modules and work flow of our implementation. The AP's probing packet is first constructed using the precoding mechanism described in §4.2. Upon receiving the probing packet, each user computes its orthogonality metric, and then starts distributed ranking. The winner feeds back its implicit CSI, which is implemented using the channel reciprocity calibration mechanism in 802.11 [18, Ch. 20.3.12.1]. The performance of MU-MIMO beamforming based on implicit CSI feedback has been validated in [7]. The winner's CSI is in turn leveraged to construct the next probing packet. After each AP selects its users to serve/null, the AP models per-user SNR (§4.4) and selects the corresponding modulation rate. MU-MIMO transmission and interference nulling start immediately afterwards.

The MU-MIMO transmission module of Kardia builds on top of 802.11 OFDM PHY. By default, it has a 20 MHz channel divided into 64 subcarriers, including 48 data subcarriers, 4 pilots, 1 DC and 11 guardband subcarriers. Modulated symbols (*e.g.*, using BPSK) on each data subcarrier are first precoded using ZFBF, mapped to different transmit antennas, and then modulated by OFDM. Standard 802.11 preambles, including STF (for synchronization), LTF (per-antenna channel estimation) and VHT-LTF (estimating the channel of precoded symbols), are prepended to the OFDM-modulated symbols. At the receiver side, correlation-based algorithms are used to detect incoming packets, estimate the CSI, and compensate for frequency offset. The resulting OFDM symbols in the packet are then demodulated and mapped into data bits. Note that the OFDM-symbol-level synchronization requirement of interference nulling is natively satisfied by the WARP platform.

Due to the processing latency and interface delay between the WARP radio and its PC host, we cannot directly reproduce MAC/-PHY level mechanisms that require instantaneous feedback/response or in-hardware computation. In particular, precoding/decoding a packet costs several hundred milliseconds in WARPLab, which is orders of magnitude longer than that in real WiFi hardware. To circumvent this problem, we build a backend module on top of WARPLab, which uses a virtual clock for time-critical operations. Kardia's key MAC components including MOP, OAS, beamforming rate adaptation, and PHY layer MU-MIMO transmission and interference nulling, are first run over the air. Then, the virtual clock replaces the WARP-induced inter-frame latency with the response time of real WiFi hardware (the 9 $\mu s$ SIFS for inter-frame time). This ensures the achievable throughput of each transmission is faithfully reproduced. The resulting network performance is thus consistent with that when Kardia is running in actual WiFi radios.

# 7. EVALUATION

In this section, we first evaluate each module of Kardia through micro-benchmark experiments based on the testbed implementation detailed in §6. We then conduct a system-level test to benchmark Kardia against the state-of-the-art MU-MIMO schemes.

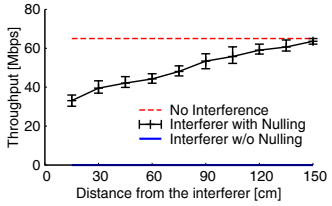


**Figure 11: The effectiveness of interference nulling. Error bars indicate max/min across 10 experiment runs.**

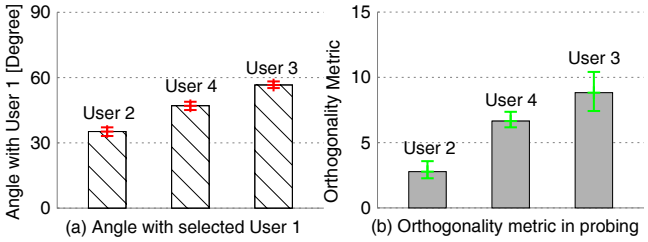


**Figure 12: MOP's probing result faithfully reflects the orthogonality between users.**

## 7.1 Micro-benchmark evaluation

**Validation of interference nulling.** To verify the efficacy of inter-cell interference nulling in our Kardia implementation, we measure the throughput of a SISO link (link distance 3 m) in the presence of an interfering transmission. The interferer nulls its transmission at the SISO receiver through zero-forcing precoding. Figure 11 shows the throughput achieved at the SISO receiver when it is separated by varying distances from the interferer. Without interference nulling, no frame can be correctly received. If interference nulling is enabled, the receiver can almost achieve its interference-free throughput when the distance separation is beyond 120 cm. Below that separation, leakage interference still exists and may halve the throughput in extreme cases. In practice, the typical distance between a user and interfering AP is larger than that separation, interference nulling should be even more effective.

**Validation of orthogonality probing in MOP.** We verify MOP's orthogonality metric design using a 4-antenna AP, which first selects a user (user1), and then probes all 3 remaining directions that are orthogonal to user1's channel. Figure 12 shows that the unselected users' performance metrics are closely correlated with their channel angles *w.r.t.* user1. The "more orthogonal" (closer to $\frac{\pi}{2}$) it is, the higher the orthogonality metric. Therefore, this metric can effectively single out users most orthogonal to existing ones.

**Optimality of DoF distribution.** In theory, Kardia's DoF distribution can achieve at least 50% of the optimal throughput (§5). We now experimentally compare Kardia with a an *oracle* that obtains a globally optimal user selection solution via an exhaustive search based on global CSI knowledge. Both Kardia and the oracle run in our testbed topology with one 4-antenna AP and 10 randomly located users. Figure 13 shows CDF of per-user throughput. We observe that Kardia's performance closely approximates that of the oracle solution, with a median loss of only 8%. Due to exponential complexity of the exhaustive search, we were unable to compute the oracle result for multi-AP case. We will evaluate this case shortly using 802.11ac as benchmark (§7.2).

**Proportional fairness control.** Recall Kardia uses a weighted 'first-user selection' scheme to approximate proportional throughput fairness (Sec 4.5). To validate its effectiveness, we compare the throughput performance of Kardia with and without such fairness control. Figure 14 plots the per-user throughput distribution of Kardia with and without proportionally fair user selection, in contrast

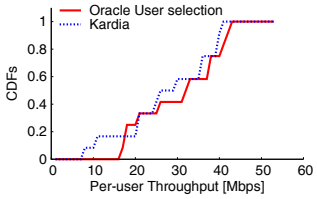


**Figure 13: MOP's user selection in contrast to an oracle.**

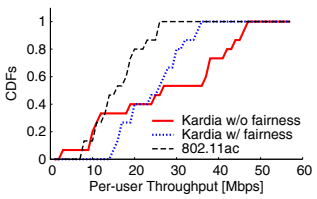


**Figure 14: Per-user throughput and fairness of Kardia.**

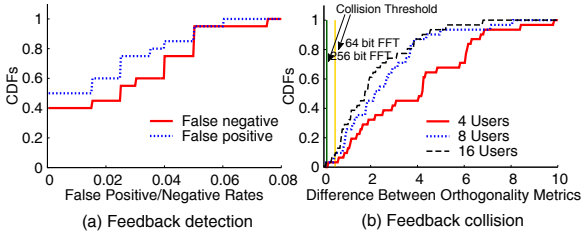


**Figure 15: Performance of distributed ranking.**

to 802.11ac. The results are obtained from 3 testbed topologies, each with the same number of AP/users as above. Observe the median throughput levels of Kardia with and without proportionally fair user selection are similar. But, the throughput variance of Kardia is smaller with fairness control enabled. It delivers substantial throughput gain over 802.11ac, for all users in the network.

**Distributed ranking.** Next we verify Kardia's frequency-domain distributed ranking. Figure 15(a) shows the CDF of false positive/negative rates in a 10-user topology running Kardia, where false positive is defined as a user wins the distributed ranking that it did not join. Recall that not all users will join each ranking due to the pruning of low orthogonality-metric users. Thus, both false positive and negative rates are very low. Consequently, the AP can accurately single out the user with highest metric.

Collisions happen when the difference of users' orthogonality metrics fall lower than the resolution of quantization mechanism (§4.2). However, Figure 15(b) shows that the probability of such collisions under various topologies is below 5% when 64-subcarriers are used. With 256 subcarriers, the resolution of orthogonality metric quantization improves and the collision in distributed ranking becomes negligible. Note that collisions on lower-index subcarriers will not affect the performance of Kardia. It only matters when both of the colliding users have the same highest orthogonality metric, and thus choose the same set of upper subcarriers.

**Accuracy of rate adaptation.** We next investigate the accuracy of the SNR estimation method (§4.4) using two 4-antenna APs. Each AP transmits to 2 users and nulls the another 2 users served by the other AP. Figure 16 shows the difference between our estimation and ground truth (measured through post-processing of received packets), distributed across 30 different real-world topologies. Observe that our SNR estimation method bears a maximum error of 1.88dB and average 0.65dB, which suffices for bit-rate adaptation with a reasonable accuracy.

**Effectiveness of OAS.** We now verify the OAS (§4.3) in a random topology with a 2-antenna AP and 10 users. Figure 17 shows that Kardia's simple subcarrier selection can improve the mean data rate of all users from about 55Mbps to 75Mbps — a 36% increase. Note that this increase comes despite the fact that not all OFDM subcarriers are utilized. Therefore, the OAS can accurately avoid poorly performing subcarriers to improvement overall throughput.

**Kardia's overhead.** The overhead in Kardia includes *(i)* the orthogonality probing frame, *(ii)* the distributed ranking and *(iii)* the implicit CSI feedback. Because the WARPLab platform cannot meet 802.11ac timing demands, we evaluate the overhead using a

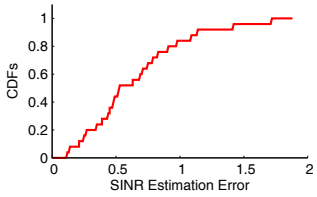


**Figure 16: Accuracy of Kardia SNR estimation.**

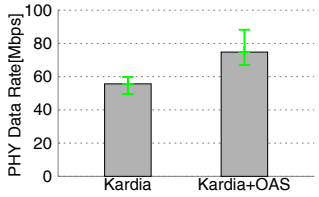


**Figure 17: Average (and std.) rate over 10 users using OAS.**

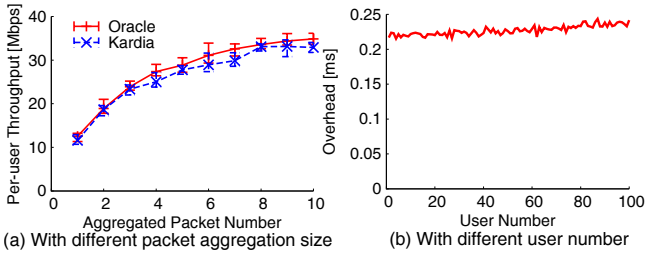


**Figure 18: Overhead analysis for Kardia.**

simulated network (§6) of 10 4-antenna APs and 100 users. We compare the overhead of Kardia with that of an oracle which has no overhead for user/AP selection, but still incurs the implicit CSI feedback overhead needed for MU-MIMO transmission.

Figure 18(a) shows that Kardia has similar performance to the oracle scheme under a variety of frame aggregation levels (with 1.5 KB per-frame). This is because *(i)* orthogonality probing only incurs short preambles, and *(ii)* each distributed ranking only occupies a short period of 17 $\mu s$ (§4.2). Note that the throughput of both Kardia and oracle scheme degrade significantly with low frame aggregation level. This is attributed to the unavoidable probing/feedback overhead in all closed-loop MU-MIMO networks.

Figure 18(b) further reveals Kardia's scalability as its overhead is almost independent of user population. We can observe that the overhead only increases slightly under much larger user number. This increase is caused by collisions in the distributed ranking. Since the collision probability remains very low under large user number, the amount of such increase is trivial.

## 7.2 System level evaluation

In this section, we integrate all Kardia components for performance evaluation in practical topologies.

**Small scale experiments.** We first test Kardia in a small scale prototype implementation of 2 APs and 6 users. Kardia is compared with (a) random nulling, where the 2 APs randomly select users to serve and null, and (b) vanilla 802.11ac, which only allows single cell transmission. We vary the location of APs and users and repeat the experiment for 30 different topologies. Figure 19 shows the CDF of total network throughput among all topologies. On average, Kardia improves the total network throughput by 54% over vanilla 802.11ac and 67% over random nulling. The performance gain here comes from both orthogonality-based user scheduling within each cell and the DoF distribution over the two cells.

**Performance in large scale networks.** The DoF distribution capability of Kardia reaps even greater throughput gains in larger and denser networks. Here, we evaluate Kardia over a larger network of 5 APs and 20 users, partitioned into 2 contention domains following §4.5. Due to a limited number of available WARP boards, we adopt a trace-based emulation approach similar to [8]. We collect channel matrices between all APs and clients in the testbed locations (Figure 21). We then emulate Kardia by replacing the WARP hardware with per-frame channel distortion following the traces.

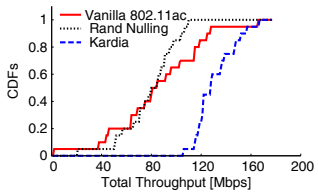


**Figure 19: Throughput in small scale network.**

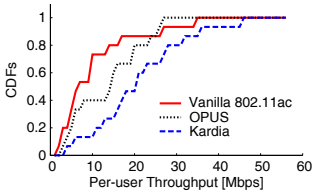


**Figure 20: Throughput in large scale network.**

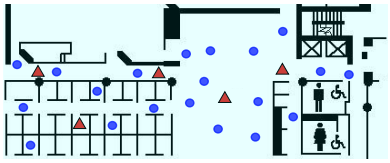

**Figure 21: The testbed topology. Blue and red markers denote the user and AP locations, respectively.**

Figure 20 compares the per-user throughput of different protocols, in which we also include OPUS [9], an orthogonality-aware user selection scheme that essentially limits itself to intra-cell DoF distribution. Experimental results reveal that *(i)* vanilla 802.11ac has the worst performance, with a total network throughput of only 152Mbps. *(ii)* OPUS outperforms vanilla 802.11ac, but its intra-cell user selection alone does not explore the full network capacity. *(iii)* Kardia achieves the highest performance, with a total throughput of 314 Mbps. The capability of DoF distribution amongst different cells allows Kardia to achieve 2.07× and 1.56× of average throughput gain over vanilla 802.11ac and OPUS, respectively.

# 8. RELATED WORK

**Multi-cell coordinated beamforming.** Coordinated beamforming has been well explored in infrastructure networks [19, 20]. A common theme lies in joint optimization of precoding matrix among APs to maximize capacity. Such optimization is solvable only for a given set of APs and users with known CSI. A recent algorithm [21] allows a MIMO node to transmit through a direction orthogonal to existing transmissions, but it requires lengthy learning phase. ProBeam [22] exploits switched-beamforming to optimize spatial reuse in WiMax. Owing to a discrete set of beam directions, the problem can be cast into a matching between the transmitters' beams and receiver set, which is further mapped to the generalized assignment problem [23]. In contrast, Kardia runs ZFBF beamforming that derives a unique set of beams for each user combination, rendering the user/AP selection problem much harder.

Recently, multi-cell coordinated beamforming in WLAN has become a research hotspot. However, existing works [24, 25] mainly focus on *multi-antenna* clients. Their design principle is to make the most of the multiple DoFs at the *client* side. In contrast, Kardia focus on the optimal utilization of the DoFs at the *AP* by employing multi-cell transmission to combat the channel hardening effect. Kardia works even for *single-antenna* clients.

**User grouping in single-cell MU-MIMO.** Current wireless standards, including 802.11ac, WiMax and LTE, can run MU-MIMO only for a given set of users irrespective of their orthogonality. Substantial theoretical work has devised iterative algorithms to combat user coupling [4–6] for single-cell MU-MIMO. These algorithms aim to compute an optimal user combination, but assuming global CSI is known to the AP. Alternative algorithms [26] use the statistical instead of full CSI. But practical wireless channels are hard to be characterized by a stationary model due to multi-path fading. Kardia is partly inspired by OPUS [9], where an AP incrementally

selects users based on their orthogonality to existing ones in its own cell. Kardia addresses a more comprehensive problem of multi-cell joint selection of users and APs, and derives an algorithm with provably good performance. Our experiments verified that it is superior over a stand-alone counterpart where each AP selects users separately (Figure 20).

**Cross-layer MU-MIMO protocols.** State-of-the-art MU-MIMO systems [2, 24, 27, 28] commonly adopt random access protocols, *i.e.*, random user selection. They are employed only to maximize the number of concurrent users, rather than the sum capacity. Such systems can be further enhanced by Kardia to combat channel hardening and thus better harness the downlink DoFs.

Network MIMO [13, 29, 30] and interference alignment [31] can increase DoF linearly with the *total* number of transmit antennas in the network (assuming channel hardening is circumvented). However, both approaches rely on *coherent* signal combination/alignment between distributed transmitters, which in turn requires carrier phase and sampling clock synchronization [31]. In contrast, although Kardia advocates multi-cell cooperation, it only requires APs to synchronize their packet transmission time, and can tolerate jitters within a CP duration ($0.8\mu s$). This has proven feasible for COTS APs with a wireline backhaul [10].

Combining MU-MIMO and DAS (Distributed Antenna System) [32] can improve spatial reuse and alleviate hidden terminals. However, this design allows multiple widely-distributed transmit antennas to access the channel independently, which severely suppresses legacy 802.11ac transmissions nearby. In contrast, Kardia will naturally regress to legacy 802.11ac mode if all APs nearby ignore the extra transmission opportunities offered by the Kardia AP.

# 9. CONCLUSION

In a distributed MU-MIMO network, channel orthogonality among users has a significant impact upon the network throughput. Improper grouping of MU-MIMO downlink users, particularly those with poor orthogonality, is detrimental to the overall network throughput. We have designed and evaluated Kardia, a novel framework that can efficiently identify the throughput-maximizing users/APs combinations within a WLAN cell and across cells. System evaluations show that Kardia improves the throughput in a large-scale network by 107% over 802.11ac. Kardia has potential to enhance a broad class of MU-MIMO protocols that group users randomly.

# Acknowledgement

This research was supported in part by the NSF under Grant CNS-1318292, CNS-1343363, CNS-1350039 and CNS-1404613.

# 10. REFERENCES

[1] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[2] 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 *Proc. of ACM MobiCom*, 2012.

[3] B. Hochwald and S. Vishwanath, "Space-time multiple access: Linear growth in the sum rate," in *in Proc. 40th Annual Allerton Conf. Communications, Control and Computing*, 2002.

[4] T. Yoo, N. Jindal, and A. Goldsmith, "Multi-Antenna Downlink Channels with Limited Feedback and User Selection," *IEEE JSAC*, vol. 25, no. 7, 2007.

[5] T. Ji, C. Zhou, S. Zhou, and Y. Yao, "Low Complex User Selection Strategies for Multi-User MIMO Downlink Scenario," in *IEEE WCNC*, 2007.

[6] M. Trivellato, F. Boccardi, and F. Tosato, "User Selection Schemes for MIMO Broadcast Channels with Limited Feedback," in *Proc. of IEEE Vehicular Technology Conference*, 2007.

[7] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong, "Argos: Practical Many-Antenna Base Stations," in *Proc. of ACM MobiCom*, 2012.

[8] X. Xie, X. Zhang, and K. Sundaresan, "Adaptive Feedback Compression for MIMO Networks," in *ACM MobiCom*, 2013.

[9] X. Xie and X. Zhang, "Scalable User Selection for MU-MIMO Networks," in *Proc. of IEEE INFOCOM*, 2014.

[10] S. Kumar, D. Cifuentes, S. Gollakota, and D. Katabi, "Bringing Cross-Layer MIMO to Today's Wireless LANs," in *Proc. of ACM SIGCOMM*, 2013.

[11] K. Tan, J. Fang, Y. Zhang, S. Chen, L. Shi, J. Zhang, and Y. Zhang, "Fine-Grained Channel Access in Wireless LAN," in *Proc. of ACM SIGCOMM*, 2010.

[12] S. Sen, R. Roy Choudhury, and S. Nelakuditi, "No time to countdown: Migrating backoff to the frequency domain," in *Proc. of ACM MobiCom*, 2011.

[13] X. Zhang, K. Sundaresan, M. A. Khojastepour, S. Rangarajan, and K. G. Shin, "NEMOx: Scalable Network MIMO for Wireless Networks," in *Proc. of ACM MobiCom*, 2013.

[14] L. Bordeaux, Y. Hamadi, and P. Kohli, *Tractability: Practical Approaches to Hard Problems*. Cambridge University Press, 2014.

[15] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák, "Maximizing a Submodular Set Function Subject to a Matroid Constraint," in *IPCO*, 2007.

[16] J. Lee, V. S. Mirrokni, V. Nagarajan, and M. Sviridenko, "Maximizing Nonmonotone Submodular Functions Under Matroid or Knapsack Constraints," *SIAM J. Discret. Math.*, vol. 23, no. 4, 2010.

[17] M. Fisher, G. Nemhauser, and L. Wolsey, "An Analysis of Approximations for Maximizing Submodular Set Functions – II," *Mathematical Programming Studies*, 1978.

[18] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," *IEEE Std. 802.11n*, 2009.

[19] H. Dahrouj and W. Yu, "Coordinated Beamforming for the Multicell Multi-Antenna Wireless System," *IEEE TWC*, vol. 9, no. 5, 2010.

[20] C.-B. Chae, I. Hwang, R. Heath, and V. Tarokh, "Interference Aware Coordinated Beamforming in a Multi-Cell System," *IEEE TWC*, vol. 11, no. 10, 2012.

[21] Y. Noam and A. Goldsmith, "Exploiting Spatial Degrees of Freedom in MIMO Cognitive Radio Systems," in *Proc. of IEEE ICC*, 2012.

[22] J. Yoon, K. Sundaresan, M. A. Khojastepour, S. Rangarajan, and S. Banerjee, "ProBeam: A Practical Multicell Beamforming System for OFDMA Small-Cell Networks," in *Proc. of ACM MobiHoc*, 2013.

[23] L. Fleischer, M. X. Goemans, V. S. Mirrokni, and M. Sviridenko, "Tight Approximation Algorithms for Maximum General Assignment Problems," in *Proc. of ACM-SIAM SODA*, 2006.

[24] K. C.-J. Lin, S. Gollakota, and D. Katabi, "Random Access Heterogeneous MIMO Networks," in *ACM SIGCOMM*, 2011.

[25] H. Yu, O. Bejarano, and L. Zhong, "Combating Inter-cell Interference in 802.11ac-based Multi-user MIMO Networks," in *Proc. of ACM MobiCom*, 2014.

[26] D. Gesbert, M. Kountouris, R. W. Heath, C.-B. Chae, and T. Salzer, "Shifting the MIMO Paradigm," *IEEE Signal Processing Magazine*, vol. 24, no. 5, 2007.

[27] 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 *Proc. of ACM MobiCom*, 2009.

[28] S. Gollakota, S. D. Perli, and D. Katabi, "Interference Alignment and Cancellation," in *Proc. of ACM SIGCOMM*, 2009.

[29] H. S. Rahul, S. Kumar, and D. Katabi, "JMB: Scaling Wireless Capacity With User Demands," in *Proc. of ACM SIGCOMM*, 2012.

[30] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, "Achieving High Data Rates in a Distributed MIMO System," in *Proc. of ACM MobiCom*, 2012.

[31] O. El Ayach, S. Peters, and J. Heath, R.W., "The Practical Challenges of Interference Alignment," *IEEE Wireless Communications*, vol. 20, no. 1, 2013.

[32] J. Xiong, K. Sundaresan, K. Jamieson, M. A. Khojastepour, and S. Rangarajan, "MIDAS: Empowering 802.11ac Networks with Multiple-Input Distributed Antenna Systems," in *CoNEXT*, 2014.