MU-MIMO-Aware AP Selection for 802.11ac Networks

Yunze Zeng*
University of California, Davis, CA, USA
zeng@ucdavis.edu

Ioannis Pefkianakis
Hewlett-Packard Labs, Palo Alto, CA, USA
ioannis.pefkianakis@hpe.com

Kyu-Han Kim
Hewlett-Packard Labs, Palo Alto, CA, USA
kyu-han.kim@hpe.com

Prasant Mohapatra
University of California, Davis, CA, USA
zeng@ucdavis.edu

ABSTRACT

Major Wi-Fi Access Point (AP) vendors worldwide seek to provide gigabit wireless connectivity, by densely deploying MU-MIMO capable APs, which can support multiple, concurrent data streams to a group of clients, connected to them. However, MU-MIMO gains can only be achieved if an AP can identify groups of clients with homogenous configurations and orthogonal wireless channels, where concurrent transmissions will not cause inter-client interference. Hence, MU-MIMO performance is fundamentally depending on how the clients are assigned to APs. Our experiments with 802.11ac commodity testbeds show that state-of-the-art client assignment algorithms are MU-MIMO oblivious and limit the MU-MIMO grouping opportunities in realistic settings. In this paper, we design and implement MAPS, an MU-MIMO-Aware AP Selection algorithm that is 802.11-compliant and can boost network's MU-MIMO throughput gains. We verified MAPS' gains over legacy designs via extensive experiments with 802.11ac commodity testbeds.

CCS CONCEPTS

• **Networks** → *Wireless access points, base stations and infrastructure*;

KEYWORDS

Wi-Fi, IEEE 802.11ac, MU-MIMO, AP selection

1 INTRODUCTION

The major Wi-Fi Access Point (AP) vendors worldwide seek to provide ubiquitous, gigabit wireless connectivity, by deploying high-density networks [1, 15], where a large number of clients and APs operate in the same RF coverage zone. A key feature for such deployments is MU-MIMO, which uses beamforming to support multiple, concurrent data streams from an AP to a group of client devices. MU-MIMO feature has already been adopted by the IEEE 802.11ac networks, to realize gigabit downlink speeds. It has been also widely perceived among the primary means to meet the speed requirements of the next generation Wi-Fi 802.11ax [6] and 5G networks [3]. However, MU-MIMO gains can be achieved only if an AP can identify groups of clients with homogenous configurations and orthogonal wireless channels, where concurrent transmissions will not cause inter-client interference. Consequently, MU-MIMO performance is fundamentally depending on how clients will be assigned to APs, which operate in the same coverage zone.

Limitations of legacy designs: State-of-the-art AP selection designs proposed by industry [2, 8] and academia [17, 29] assign clients to the strongest signal (RSSI) AP. Interestingly, our experiments with commodity MU-MIMO 802.11ac testbeds show that they yield more than 50% lower network throughput compared to the optimal client assignment. We have identified three root causes for such poor performance. (*a*) Legacy designs are MU-MIMO oblivious and may assign clients with correlated channels to the same AP. Grouping clients with correlated channels results in high inter-client interference, which forces the AP to operate in SU-MIMO mode (serving one client at a time). (*b*) Even when clients communicate over orthogonal channels thanks to rich multipath environment, legacy designs may still limit MU-MIMO grouping opportunities, by assigning clients with heterogeneous bandwidth configurations to the same AP. Specifically, clients which operate on different 802.11ac bandwidth options cannot be grouped together. This constraint is attributed to AP's capability, which can only transmit on a single center frequency and bandwidth at a time. Moreover, the AP cannot always use the highest bandwidth option for transmitting to all the clients in a group, due to clients' different interference profiles. (*c*) Finally, the widely adopted approach to assign clients to the least "loaded" AP, can limit MU-MIMO gains, even for clients which operate on orthogonal channels and homogenous bandwidths. This is because more loaded APs may offer more grouping opportunities.

Design challenges: The design of MU-MIMO-aware AP selection, which addresses the above limitations, poses significant challenges. First, a key design challenge is to identify clients with orthogonal channels and assign them to the same AP. A naive approach would be to associate each client to all APs in its range, and use explicit beamforming feedback [19, 21, 26] to estimate channel correlation. However, such an approach requires excessive handoffs in high-density Wi-Fi deployments, and leads to poor performance. Hence, a new, low-overhead approach is required for profiling the multipath environment and estimating channel correlations. Moreover, MU-MIMO-aware AP selection should be able to capture clients' (and APs') bandwidth profiles, which dynamically change in time, due to interferences. Then, it needs to estimate how clients' bandwidth profiles will affect their grouping opportunities at an AP. Finally,

*Yunze Zeng completed this work during his internship at Hewlett-Packard Labs.

Mobihoc '17, July 10-14, 2017, Chennai, India

ACM ISBN 978-1-4503-4912-3/17/07…$15.00
https://doi.org/10.1145/3084041.3084057

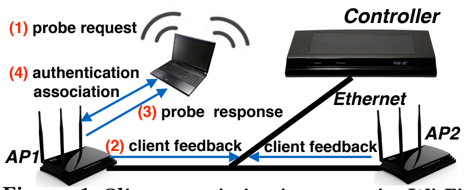

Figure 1: Client association in enterprise Wi-Fi.

AP selection must balance the load among APs, without limiting their MU-MIMO gains, which are often conflicting objectives.

MAPS design: In this paper, we propose a new *M*u-mimo-*A*ware *AP S*election (*MAPS*) design for 802.11ac networks, which addresses the aforementioned challenges. MAPS leverages a NULL frame probing scheme to collect CSI (Channel State Information) feedback from clients, without requiring them to associate with APs. CSI samples measured at the AP-side can capture the multipath characteristics of the environment, and can be used as a proxy for clients' channel correlation, as shown by our experiments. MAPS first sanitizes CSI samples by removing the RF-hardware triggered amplitude deviations, using local regression smoothing filters. It then constructs a CSI profile that differentiates between persistent and transient multipath. The CSI profiler applies a correlation metric among back-to-back CSIs, which captures dominant multipath changes, and at the same time remains robust to RF hardware-triggered CSI phase shifts. Using the CSI profile, MAPS can estimate the SINR (Signal-to-Interference-plus-Noise Ratio) and hence the PHY rate of a client as a part of an MU-MIMO group. MAPS' SINR approximation error is typically small (<2 dB) compared to SINR estimation using explicit client's channel feedback.

MAPS introduces a novel client assignment model, which leverages clients' SINR, traffic and interference profiles to infer the effective MU-MIMO throughput of a client, at each AP. Our model can balance between MU-MIMO gains and AP load, by considering the Wi-Fi channel busy time, and the airtime to be allocated to a client, at each AP. MAPS will then assign clients to APs which maximize their throughputs. Since optimal client assignment is an NP-Hard problem, we propose a low-overhead heuristic algorithm, which performs close to optimal, as shown by our experiments.

We evaluate MAPS' performance gains over legacy designs using testbed experiments with 802.11ac commodity APs and MU-MIMO-capable smartphones. Our results show that MAPS outperforms legacy designs in 90% of the settings, with network throughput gains greater than 50%. In the most (~85%) of our experiments, MAPS performs the same as the optimal, best-throughput client assignment. Our simulations using traces from Wi-Fi enterprise networks verify MAPS gains in large scale topologies, where more than 50 clients are connected to an AP.

Contributions: In summary, our main contributions are:

(1) We conduct a measurement study with commodity 802.11ac MU-MIMO testbeds, and identify the limitations of legacy AP selection designs (Sec. 3). To the best of our knowledge, this is the first work that studies AP selection in MU-MUMO networks, using commodity 802.11ac testbeds.

(2) We design MAPS, a practical, 802.11-compliant system, which can boost MU-MIMO gains by appropriately assigning clients to APs. (Sec. 4).

(3) We implement MAPS in 802.11ac commodity hardware (Sec. 5), and evaluate its performance in multiple network settings, using 802.11ac APs and smartphone devices (Sec. 6).

2 BACKGROUND

2.1 IEEE 802.11ac Background

The key differentiator of 802.11ac over its predecessors is the MU-MIMO feature, which uses beamforming to support concurrent downlink data streams from an AP to a group of clients. An 802.11ac AP can support MU-MIMO beamforming, by using a sounding protocol [10] to collect *VHT Compressed Beamforming Feedback (CBF)* from wireless clients. The CBF is represented by V, which is a steering matrix that specifies how AP should decorrelate transmitted data to multiple clients. A client calculates V by applying Singular Value Decomposition (SVD) on H as $H = UDV^H$. Here, H (or CSI) is the channel matrix measured at the client's side from sounding packet. Then, an AP uses V to precode the transmission data. Apart from CBF, 802.11ac clients provide SNR (Signal to Noise Ratio) feedback to AP. An 802.11ac AP selects a set of clients to transmit data concurrently through a *client selection* algorithm that precedes sounding and beamforming. Client selection algorithm is vendor-specific and unspecified by the 802.11ac standard. 802.11ac supports 20, 40, 80 MHz channel bandwidths, and an optional 160 MHz bandwidth. An 802.11ac device can use a 20 MHz sub-channel only if it is not occupied by another transmission in its vicinity. An 802.11ac AP can negotiate communication at higher channel widths through an RTS/CTS handshake protocol [10, 27]. Interestingly, *only clients with the same channel bandwidth configuration can be grouped together in the same MU-MIMO group*. This is because an AP can only transmit using a single center frequency and bandwidth at a time. Moreover, clients of different channel bandwidth may have different interference profiles, and hence the AP cannot transmit data to all of them using the highest channel bandwidth.

2.2 Enterprise Wi-Fi Networks

In this paper, we focus on enterprise Wi-Fi networks where multiple APs operating in infrastructure mode, provide wireless connectivity in large buildings. APs are typically connected through Ethernet to controllers, as shown in Figure 1. Controllers support network management, and enterprise applications' security. They also assign clients to APs (i.e., AP selection) and initiate clients' handoffs. Specifically, prior to association, a client C scans Wi-Fi channels for APs. In *passive* scanning mode, C listens for beacons. In *active* mode, C broadcasts probe requests. The APs which receive these requests send information about C (e.g., RSSI) to the controller, which assigns C to an AP. The selected AP sends a probe response to C, which initiates the authentication and association process. If the controller decides to handoff C from AP1 to AP2, then AP1 sends a disassociate frame to C, and C follows the process shown in Figure 1.

3 A MEASUREMENT STUDY

In this section, we show the limitations of legacy AP selection approaches in MU-MIMO networks, by conducting experiments with 802.11ac commodity wireless testbeds.

3.1 Platform and Methodology

Our experiments use commodity 802.11ac APs, equipped with a Qualcomm Beeliner 4×4 MU-MIMO-capable 802.11ac 5 GHz radio.

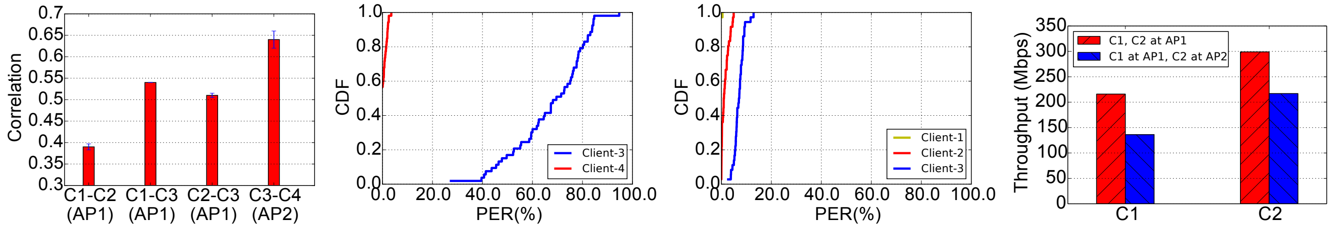

(a) Clients' V correlation. **(b) C3 and C4 at AP2.** **(c) C1,C2 and C3 at AP1.** **(d) Performance of legacy load balancing in MU-MIMO.**

Figure 2: a) V correlation, b, c) Packet-Error-Rate, d) throughput for clients in our case study settings.

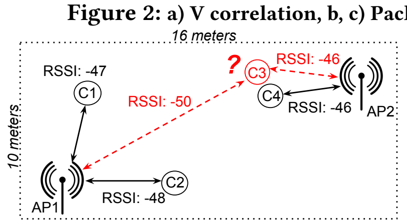

Figure 3: Network topology for the case study setting.

The 802.11ac radio supports up to 80 MHz channel bandwidth and up to 256-QAM modulation level, with 1733.3 Mbps peak PHY rate. It has 4 antennas, but supports up to 3 data streams (clients) in MU-MIMO mode. MU-MIMO client selection along with other core MAC-layer functionalities (e.g., PHY rate and bandwidth adaptation) are implemented in the AP's firmware, and the source code is available for our implementation. Our experiments use Xiaomi Mi 4i smartphones as clients, which have an 802.11ac wave-2 chipset, with one receiving antenna.

For our experiments, we have modified the firmware of our APs to collect per-client wireless feedback such as: a) PHY rates (i.e., MCS, spatial stream, channel bandwidth), b) sounding feedback statistics (i.e., V matrix, per-subcarrier SNR - cf. Sec. 2.2), c) frame error rate, and d) CSI measured at the AP side, from the received frames. We conduct experiments in enterprise and university campus settings.

3.2 Correlated Wireless Channels

We next show the limitations of legacy AP selection algorithms in typical 802.11ac deployments, with testbed experiments. In our case study setting, clients C1 and C2 are connected to the strongest signal (RSSI) AP1 and C4 to AP2, as shown in Figure 3. All clients operate at 80 MHz. C1 and C2 form an MU-MIMO group and achieve 378 Mbps aggregate downlink UDP throughput compared to only 296 Mbps, when they operate at Single-User MIMO (SU-MIMO) (where one client is served at a time). When a new client C3 wants to join the network, legacy (e.g., RSSI-based) designs will assign C3 to the highest RSSI AP2. However, is this the best-throughput client assignment?

Our results show that clients C3 and C4 perform poorly in MU-MIMO mode, at AP2. Specifically, the aggregate downlink throughput at AP2 is 178.8% higher when C3 and C4 operate in SU-MIMO, compared to forming an MU-MIMO group. Hence, C3 cannot leverage MU-MIMO gains by connecting to the highest RSSI AP2. Interestingly, the aggregate (over AP1 and AP2) network throughput is from 50 to 230 Mbps higher when C3 connects to the lower RSSI AP1 (assuming AP1, AP2 operate on orthogonal channels). Particularly, the best-throughput setting is observed when C1, C2 and C3 form an MU-MIMO group at AP1, and C4 operates in SU-MIMO at AP2. The aggregate and per-client throughput for the above settings is summarized in Table 1.

Client C3 cannot leverage MU-MIMO gains at AP2, because its wireless channel is highly correlated with C4, leading to inter-client interference. Channel correlation (and hence interference) between clients i, j at subcarrier s, can be estimated by the V matrix correlation [21] as:

$$\rho(i,j) = \frac{\sum_s ||V_i(s)V_j^H(s)||}{\sqrt{\sum_s ||V_i(s)||^2}\sqrt{\sum_s ||V_j(s)||^2}} \tag{1}$$

Figure 2a shows that the V correlation between C3 and C4 ($\rho(C3, C4)$) is 64% higher than that of clients C1 and C2, which gives the best MU-MIMO performance. Highly correlated channels result in inter-client interference and consequently to high Packet-Error-Rate (PER). Figure 2b shows that the median PER for C3 when grouped with C4 is approximately 70%, and the maximum PER exceeds 90%. On the other hand, the channel correlation among the clients of the MU-MIMO group {C1, C2, C3} at AP1, is lower than {C3,C4}, as shown in Figure 2a. Due to lower inter-client interference, the PER for C3 at AP1 is lower than 10% for 90% of the samples, as shown in Figure 2c. Hence, the best-throughput AP selection algorithm needs to assign C3 to the lower RSSI AP1, in order to leverage the MU-MIMO gains.

Summary: *RSSI-based AP selection designs are oblivious to MU-MIMO feature, and assign clients to APs, without considering the channel correlation among clients connected to the same AP. Correlated channels lead to high inter-client interference and low throughput, in MU-MIMO settings.*

Table 1: Per-client and aggregate (across APs) downlink UDP throughput (Mbps), when C3 connects to AP1, AP2.

Setting	C1	C2	C3	C4	Aggr.
AP1(MU),AP2(MU),C3-AP2	186 ± 5	192 ± 6	9 ± 2	92 ± 21	479
AP1(MU),AP2(SU),C3-AP2	186 ± 5	192 ± 6	147 ± 6	134 ± 5	659
AP1(MU),AP2(SU),C3-AP1	146 ± 2	148 ± 2	145 ± 2	269 ± 5	708

3.3 Heterogeneous Bandwidth Clients

We next evaluate a two-AP topology setting (similar to Fig. 3), where C1 and C2 are connected to AP1 and AP2 respectively. C1 operates at 80 MHz and C2 at 40 MHz (due to the interference from neighboring networks). Let's consider a new client C3 operating at 80 MHz, whose RSSI is -50 dbm and -46 dbm from AP1 and AP2, respectively. RSSI-based designs will assign C3 to AP2, without considering that different bandwidth clients cannot be grouped together (cf. Sec. 2.2). Consequently, C2 and C3 will operate at SU-MIMO. On the other hand, C3 could form an MU-MIMO group with C1 at AP1, increasing by 40 Mbps the aggregate network throughput (cf. rows 1, 3 of Tab. 2). Even when AP's bandwidth adaptation algorithm allows for C2 and C3 to form an MU-MIMO group at 40MHz, at AP2[1], assigning C3 to the lower RSSI AP1 still gives better network throughput, as shown in rows 2, 3 of Table 2.

Assigning clients to APs with high MU-MIMO gains can also improve fairness. For example the Jain Fairness Index[12] of the

[1]Commodity 802.11ac APs do not adjust channel bandwidth, to increase MU-MIMO grouping opportunities.

network when C3 operates in SU-MIMO at AP2 is 0.82 (1 implies perfect fairness), while it increases to 0.99, when C3 connects to AP1.

Summary: *RSSI-based AP selection designs can limit MU-MIMO grouping opportunities, by assigning clients with heterogeneous bandwidths to the same AP.*

Table 2: Per-client and aggregate downlink UDP throughput (Mbps) and fairness, for heterogeneous width clients.

Setting	C1	C2	C3	Aggr.	Jain Idx
AP1(SU),AP2(SU) C3-AP2(80MHz)	269 ± 5	80 ± 4	151 ± 5	500	0.82
AP1(SU),AP2(MU) C3-AP2(40MHz)	269 ± 5	127 ± 10	125 ± 10	521	0.87
AP1(MU),AP2(SU) C3-AP1(80MHz)	186 ± 5	161 ± 3	192 ± 6	539	0.99

3.4 Load Balancing

We finally evaluate a topology of two APs in the same vicinity, operating on the same wireless channel. This is a realistic setting in dense AP deployments, when there are not enough non-overlapping channels for adjacent APs. The approach of assigning the same channel to adjacent APs has been also used by industry [1] for better interference management and faster inter-AP handoff. In our setting, client C1 operating at 80 MHz is connected to AP1, while AP2 does not serve any client. We assume that AP1 fully utilizes the wireless channel capacity to serve C1's traffic. Let's now consider a new client C2 operating at 80 MHz, whose RSSI from AP1 and AP2 is the same. Existing designs [17, 29] will assign C2 to AP2, to balance the load among APs. However, such assignment is suboptimal in terms of throughput, as shown in Figure 2d. Clients C1 and C2 operate in MU-MIMO when both are connected to AP1, and achieve 515 Mbps aggregate throughput. However, when connected to different APs, they share the wireless medium, and hence achieve 162 Mbps lower aggregate throughput.

Summary: *Legacy load balancing designs can reduce MU-MIMO grouping opportunities and hence throughput performance, by assigning clients to the least loaded AP.*

4 DESIGN

In this section we present *MAPS* (*M*u-mimo-*A*ware *AP S*election), an 802.11-compliant system, which can boost MU-MIMO gains by appropriately assigning clients to APs. MAPS seeks to increase MU-MIMO grouping opportunities by setting three key design goals: (*a*) to accurately identify clients with uncorrelated channels at low overhead, and assign them to the same AP, (*b*) to assign clients with homogenous bandwidth settings to the same AP, by monitoring their interference profiles, (*c*) to allow for client assignment to more "loaded" APs, if they can form high-throughput MU-MIMO groups. MAPS is a practical, lightweight design, which can be implemented in commodity 802.11ac hardware, without client-side modifications.

MAPS architecture: An overview of MAPS is shown in Figure 4. MAPS takes an implicit feedback approach to identify clients with uncorrelated channels, without requiring them to be associated to an AP. It leverages a NULL data probing scheme to collect CSI samples at APs for each client in their vicinity. The compressed CSI samples are sent to the controller, which upon sanitizing them, it constructs the dominant multipath profile of a client, at each AP. The controller uses a client's CSI profile to estimate its performance (i.e., PHY rate) as a member of an MU-MIMO group. A MAPS' AP further maintains client's bandwidth and traffic profiles, which along with AP's Wi-Fi channel busy time and load, are used to

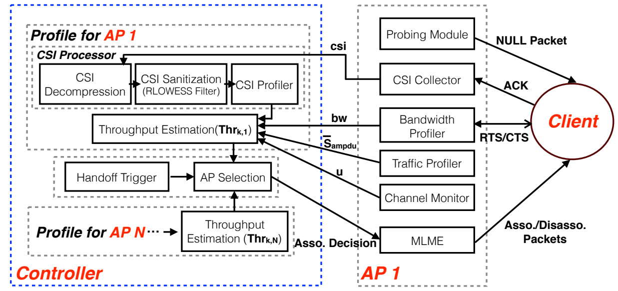

Figure 4: MAPS architecture.

estimate the throughput of a client at an AP. When a new client wants to join the network, or a client's handoff is required, MAPS' controller identifies the APs in client's range, and assigns it to the AP which maximizes its throughput. Then, it sends its decision to the selected AP, which uses the MAC Sublayer Management Entity (MLME), to associate the client to itself. We next elaborate on MAPS' building blocks.

4.1 MU-MIMO Performance Inference

A key challenge for MAPS is to identify clients with uncorrelated channels, and map them to the same AP. A naive approach would be to periodically associate each client to all APs in its range, and collect CBF to estimate clients' channel correlation. However, such approach requires frequent handoffs and long associations with low-throughput APs. Given the dense Wi-Fi deployments with multiple APs in a client's range and the long handoff times (order of seconds), explicit feedback approach will perform very poorly. Hence, MAPS takes an implicit feedback approach.

4.1.1 Leveraging Implicit Feedback. MAPS uses implicit, AP-side, CSI feedback to identify MU-MIMO groups of clients with uncorrelated channels.Intuitively, since the "physical" wireless channel is reciprocal [20], CSI measured at the AP can capture the correlation among clients' channels. However, leveraging AP-side CSI poses several challenges. How can an AP collect CSIs from clients not associated with it? Is it possible to filter Wi-Fi RF-hardware triggered noises of the collected CSI samples? How to construct a CSI profile that captures both persistent and transient multipath characteristics of the environment? We next describe MAPS' approach to such challenges.

CSI collection: MAPS leverages a NULL data probing scheme to collect CSIs from the clients in AP's vicinity. Specifically, an AP transmits NULL frames, and estimates the CSI from the ACKs sent by the client. A CSI sample is a $N_t \times N_r$ matrix of complex numbers reported per OFDM subcarrier, where N_t and N_r are the number of antennas at the AP and client. A client will respond to a NULL frame received by an AP, even if it is not connected to this AP, as verified by our experiments.

NULL probing is a low overhead CSI collection scheme. Particularly, a NULL frame is only tens of bytes (depending on the 802.11 family). Hence, its transmission time is only 0.85 microseconds, for the smartphones used in our testbed. This overhead is negligible considering that MAPS only periodically (every 100ms) transmits NULL frames.

CSI sanitization: Our measurements show that CSIs reported by commodity APs can be noisy. Such noise is attributed to transmission power changes, rate adaptation, internal CSI reference level

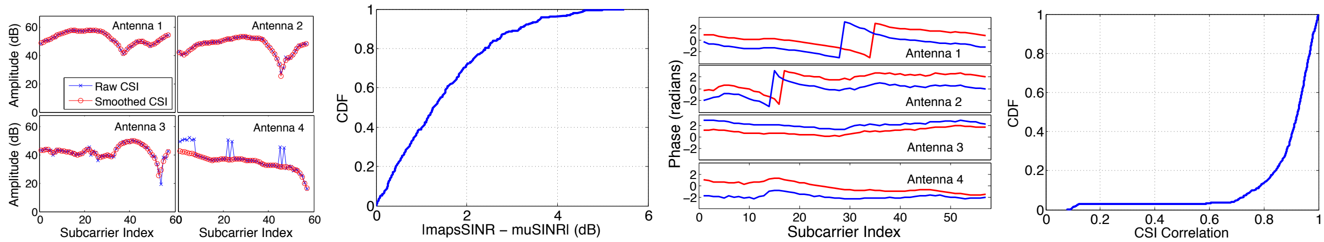

(a) Robust LOWESS filter on CSI samples. **(b) Implicit SINR vs. muSINR in multiple settings.** **(c) Phase shift for 2 back-to-back frames.** **(d) CSI correlation for back-to-back frames.**

Figure 5: MAPS' CSI profiling.

changes [22]. For example, Figure 5a shows that noise spikes can exceed 10 dB. MAPS applies a robust *LOWESS* (Locally Weighted Scatterplot Smoothing) filter [5], which performs local regression with weighted least squares to smoothen outliers. Non-parametric smoothers like LOWESS are appropriate for CSIs, since they do not assume that the data fit some distribution shape. Figure 5a shows that such filter removes noise from CSIs.

The smoothed CSI is used to estimate the SINR of a client k operating in an MU-MIMO group of K clients as [21]:

$$SINR = \frac{\underbrace{\frac{1}{K}||D_k||^2}_{signal\_power}}{\underbrace{N}_{noise\_floor} + \underbrace{\frac{1}{K}||D_k||^2 \Sigma_{j \neq k} ||V_k^H V_j||^2}_{interference}} \quad (2)$$

MAPS estimates V and D by applying SVD on CSI (cf. Sec. 2.2). It computes the noise N using EVM (Error Vector Magnitude) feedback, provided by AP's firmware, for every received frame across all subcarriers. Finally, it calibrates D_k to account for the transmit power difference between the client and AP. Specifically, it multiplies the factor $||D_k||^2$ with $10^{\frac{P_{AP}-P_{client}}{10}}$, which is the transmit power difference (dBm) at AP and client sides. P_{client} is available at the AP through 802.11 Event Report frames. Notice that the SINR metric can be estimated per OFDM subcarrier. MAPS computes an *effective SINR* across all subcarriers using the approach proposed in [7], which has been shown to be robust in frequency-selective fading environment.

SINR accuracy: We evaluate our SINR metric's accuracy, by comparing it with muSINR [21], which uses explicit, receiver-side V and D feedback (CBF). Figure 5b shows the distribution of the absolute difference between the two SINR metrics, from multiple experimental settings. We observe that for 70% of the cases, the SINR estimation error is less than 2 dB. This error will not lead to erroneous PHY rate estimation most of the times (cf. Tab. 22-25 in [10]), and hence it does not affect MAPS' ability to infer client's throughput. Since MAPS' SINR estimation is not used for core functions such as rate adaptation, estimation error outliers will not significantly impact MAPS' performance, as shown by our evaluation results (cf. Sec. 6).

4.1.2 Identifying Dominant Multipaths. MAPS' operations are triggered at coarser time scales (sec.) compared to other CSI-based algorithms, such as MU-MIMO grouping (msec.), to avoid excessive handoffs. Hence, instead of simply maintaining the latest CSI, MAPS needs to construct a CSI profile which captures both persistent and transient multipath characteristics of the environment.

Constructing a CSI profile is a challenging task. First, storing and processing all the measured CSIs (collected at msec. scales) is a big overhead even for a Wi-Fi controller. Hence, MAPS needs to consider only CSIs that capture multipath changes. However, capturing such dynamics is not trivial. Even if the multipath characteristics of the environment remain the same, the phase of back-to-back CSI samples may vary due to Wi-Fi RF hardware characteristics [20, 22].A CSI profiler should be robust to such variations and capture only significant multipath changes.

Interestingly, different from related studies [20], our experiments show that such phase variations are not random. For example, Figure 5c shows the phase curves of two back-to-back frames, in a controlled environment with no object movement. We observe an almost constant phase shift for all subcarriers (y-axis), and a small shift in phase curve across frequency domain (x-axis), for antennas 1, 2. However, we observe similar shapes of phase curves. Since the hardware-triggered phase shifts do not change the shape of the phase curves of back-to-back frames, we expect that their correlation will be high. Figure 5d shows the distribution of the correlation factor ρ (eq. (1)) for CSIs, collected from multiple settings in a time window of 0.5 msecs, of stable multipath environment[2]. We observe that for 80% of the cases, ρ is equal or greater than 0.85 ($\rho = 1$ for same CSI samples).

While CSI correlation metric is robust to hardware-triggered phase shifts, it can also capture changes in multipath environment. We illustrate our point by studying the spatial distribution of CSIs (in polar coordinate system) in cases of stable (Fig. 6) and dynamic (Fig. 7) multipath environments. For stable multipath, back-to-back CSIs overlap in space, which is reflected in their CSI correlation. For dynamic multipath, the CSIs' main lobes do not overlap (cf. Fig. 7), as indicated by their lower correlation.

MAPS leverages the CSI correlation metric to maintain L dominant multipaths (i.e., CSIs). For each new CSI i, MAPS estimates its correlation $\rho(i, j)$ with each CSI j, of current CSI profile. If the maximum correlation with a CSI j is greater than a threshold ($max_{j \in L}\{\rho(i, j)\} > R$), then MAPS replaces CSI j, with i and increases a counter csi_j (j is an index of the CSI profile)[3]. Otherwise, the new CSI is stored in a new entry of the profile. If the CSI profile is full (with $|L|$ CSIs), MAPS will either replace an existing CSI j with the new CSI, if $csi_j = 1$, or it will discard i. MAPS periodically resets CSI profiles to allow for new multipaths.

MAPS needs to estimate clients' SINRs for all the MU-MIMO group assignments, to select the best AP (cf. Sec. 4.3). However, a client's SINR will change for each CSI in its profile. Processing all CSIs for all groups, results in significant processing overhead. MAPS amortizes such overhead, using the counter csi_i, which reflects the "dominance" of a multipath. It selects as client's dominant multipath, its CSI i with a probability $P_{csi_i} = csi_i / \sum_{j=1}^{|L|} csi_j$. Then, it uses the

[2] We compute ρ by using CSI H instead of V in eq. (1).
[3] R can be set to 0.85, from our experiments in Fig. 5d.

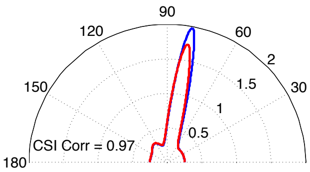

Figure 6: CSIs for stable multipath.

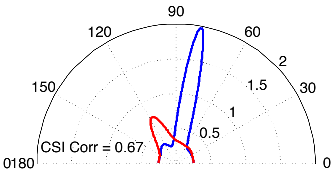

Figure 7: CSIs for dynamic environment.

dominant CSI for computing the SINR. Client's SINR along with its bandwidth and traffic profiles are used for estimating client's throughput at an AP, as we discuss next.

4.2 Client and AP Profiling

MU-MIMO grouping opportunities depend on clients' bandwidth configuration (cf. Sec. 3.3) and traffic profiles.

Bandwidth profile: MAPS monitors the interference profile of each AP and client, and connects homogenous bandwidth clients to the same AP. Particularly, it maintains for each device, the number of transmissions or receptions that could use the bandwidth option $bw \in \{20, 40, 80, 160\}$. MAPS estimates bw, by monitoring which 20 MHz sub-channels are occupied through the RTS/CTS handshake process. Since, a device's interference profile (and hence bandwidth) may change at runtime, MAPS selects bw_k for a device k, to be the bandwidth option with the highest probability. The probability of a bandwidth i is estimated as:

$$P_{bw^i} = \frac{\#Packets\_bw^i}{\sum_{j \in \{20,40,80,160\}} \#Packets\_bw^j} \tag{3}$$

MAPS maintains different bandwidth profiles for different channels that APs may operate on. It periodically resets bandwidth profiles, to account for new interference dynamics. Finally, it computes the bandwidth configuration of a client k at an AP α as: $bw_{k,\alpha} = min\{bw_k, bw_\alpha\}$.

Traffic profile: MAPS monitors the size of 802.11ac aggregated frames (i.e., A-MPDU) to capture traffic dynamics. It maintains a moving average of a client's A-MPDU size as:

$$\overline{S}_{ampdu} = (1-\beta) \cdot \overline{S}_{ampdu} + \beta \cdot S_{ampdu} \tag{4}$$

where $\beta = 1/8$ in our implementation. It periodically ages the traffic profile as: $\overline{S}_{ampdu} = (1-\beta) \cdot \overline{S}_{ampdu}$, to consider client's idle time. Typically, small $\overline{S}_{ampdu}$ implies low traffic.

MAPS also maintains the wireless channel utilization u_α and the number of clients connected to each AP α. It captures the AP's channel busy time with a factor $u_\alpha \in [0, 1]$, which is the fraction of free transmission cycle opportunities within the last 10 beacon intervals. We next show how such profiles are used to estimate clients' throughput.

4.3 AP Selection Model

MAPS introduces a novel client assignment model, which can boost MU-MIMO gains. It first leverages clients' profiles to estimate the effective throughput performance of a client, at each AP. Then, it selects the AP which can maximize a client's throughput. We next elaborate on our model.

4.3.1 Throughput Model. MAPS estimates the throughput of a client k at an AP α, considering: (1) client's PHY rate $r_{k,\alpha}$, when it operates in an MU-MIMO group at α, (2) client's traffic $\overline{S}_{ampdu,k}$, (3) 802.11 protocol overheads (t_o), (4) channel busy time u_α, and (5) airtime allocated to client $w_{k,\alpha}$. Specifically, it is:

$$Thr_{k,\alpha} = w_{k,\alpha} \cdot u_\alpha \cdot \frac{\overline{S}_{ampdu,k}}{t_d + t_o} \tag{5}$$

The factor $u_\alpha \in [0, 1]$ captures the WiFi channel busy time. The amount of client's data $\overline{S}_{ampdu,k}$ is computed based on equation (4). The time consumed by 802.11 protocol overheads (e.g., sounding, ACK) is t_o. Finally, the data transmission time is modeled as $t_d = t_{plcp} + \overline{S}_{ampdu,k}/r_{k,\alpha}$, where t_{plcp} is the PLCP preamble transmission time, and $\overline{S}_{ampdu,k}/r_{k,\alpha}$ is the frame transmission time, where $r_{k,\alpha}$ is the PHY rate.

PHY rate $r_{k,\alpha}$: For a given MU-MIMO group, MAPS first estimates client's SINR from equation (2). Then, it uses the 802.11ac rate tables [10] to map the SINR to a PHY rate $r_{k,\alpha}$ (i.e., MCS, spatial streams). However, a client can form multiple different MU-MIMO groups at each AP. These groups may change in time, depending on clients' channel correlation characteristics, bandwidth and traffic profiles. Hence, to estimate $r_{k,\alpha}$, MAPS first needs to estimate which are likely going to be k's MU-MIMO groups at AP α.

In practice, only the best-throughput MU-MIMO groups of traffic active clients are used for transmission. To this end, MAPS considers only the active clients (i.e., $\overline{S}_{ampdu} > 0$) as candidates for grouping. Then, it computes the set of MU-MIMO groups G_α that can be formed by the active clients associated to an AP α, subject to the bandwidth constraint. Here, a group $g \in G_\alpha$ is a set of clients of cardinality of at least one ($|g| \geq 1$). From all the possible MU-MIMO group combinations G_α, an AP serves only the best-throughput MU-MIMO groups, while ensuring that all clients will be fairly served. Given a set of clients K_α at AP α, a fair allocation will assign k at least $1/|K_\alpha|$ of the airtime. MAPS satisfies the above constraints, by first ordering the set G_α, in decreasing MU-MIMO group throughput order[4]. Then, starting from the highest throughput group $g \in G_\alpha$, MAPS adds g to a new set G'_α, subject to two constraints:

(a) Group g cannot be a proper subset (or superset) of existing groups in G'_α: $\forall g_i, g_j \in G'_\alpha$ then $g_i \not\subset g_j$.

(b) Every client associated to the AP α must belong to at least one group in G'_α: $\forall k \in K_\alpha$ then $k \in g$ for $g \in G'_\alpha$.

The above algorithm selects the best-throughput sets, while ensuring that all clients will be served at least once. Hence, the set $G'_\alpha \subset G_\alpha$ includes the MU-MIMO groups which will be likely served by the AP. Note that MAPS can leverage any of the techniques proposed in the literature [19, 21] to identify G'_α, at low computational cost.

Now, let's assume that $G'_{k,\alpha} \subseteq G'_\alpha$ is the subset of groups which include client k. Then, we estimate $r_{k,\alpha}$ as the average client's PHY rate when it is part of the groups $g \in G'_{k,\alpha}$: $r_{k,\alpha} = \frac{1}{|G'_{k,\alpha}|} \cdot \sum_{g \in G'_{k,\alpha}} r^g_{k,\alpha}$.

Allocated airtime $w_{k,\alpha}$: MAPS uses a weight $w_{k,\alpha}$ to capture the airtime to be allocated to client k, at AP α. Factor $w_{k,\alpha}$ depends on number of groups which are likely to be served by AP α (i.e., $|G'_\alpha|$), or by other APs in α's vicinity, operating on the same channel. Specifically, the airtime to be allocated to a client k at α is the cardinality of subset of groups that include k ($G'_{k,\alpha} \subset G'_\alpha$), to the total number of groups. Hence, we set $w_{k,\alpha}$ to be equal to

[4] Group throughput is estimated from eq. (5), by considering the group's total transmitted data and transmission time.

$|G'_{k,\alpha}|/|G'_{\alpha}|$. For example, in the case study scenario of Figure 3, $G'_{AP1} = \{\{C1, C2, C3\}\}$, $G'_{AP2} = \{\{C4\}\}$ when C3 connects to AP1, and $G'_{AP1} = \{\{C1, C2\}\}$, $G'_{AP2} = \{\{C3\}, \{C4\}\}$ when C3 connects to AP2. Hence, C3 will get more airtime at AP1 ($w_{C3,AP1} = 1$), compared to AP2 ($w_{C3,AP2} = 1/2$). Note that, if a set of APs A in the same vicinity with α operate on the same channel, then $w_{k,\alpha} = |G'_{k,\alpha}|/\sum_{\alpha' \in A} |G'_{\alpha'}|$.

In conclusion, MAPS' throughput model can capture inter-client interference and client's bandwidth profile with the rate factor r_k. It captures Wi-Fi channel utilization u_α, and prevents client assignment to APs with congested channels. Finally, it captures the "load" w_α at each AP (or APs in range, on the same channel), which is required for *load balancing*. Our AP platform maintains all the per-client state, which is required for throughput estimation.

4.3.2 Client Assignment Model. MAPS' objective is to determine the proper client-AP association set K_α, $\forall \alpha \in A$, such that the sum-throughput of all clients can be maximized. This optimization problem can be formulated as:

$$\begin{aligned} I^* = \underset{I}{\operatorname{argmax}} \sum_{k \in K} \sum_{\alpha \in A} I_{k,\alpha} Thr_{k,\alpha} \quad & \text{subject to} \\ \sum_{\alpha \in A} I_{k,\alpha} \leq 1, \forall k \in K & \\ I_{k,\alpha} \in \{0, 1\}, \forall k \in K, \forall \alpha \in A & \end{aligned} \tag{6}$$

where the binary variable $I_{k,\alpha}$ indicates whether a client k associates with an AP α,$\forall \alpha \in A$. Such constraint ensures that each client associates with at most one AP.

We prove that our problem is *NP-Hard*, by reducing a simple instance of it, to the maximum independent set problem. We describe the reduction in our technical report [28]. Hence, we propose a heuristic algorithm that seeks to maximize the aggregate clients' throughput (eq. (6)), while satisfying all the constraints. The algorithm operates as follows.

Profiling: MAPS periodically updates each client's and AP's profile. CSI profiles are updated every 100 ms, upon NULL frame transmission. Bandwidth profiles are updated upon RTS/CTS handshake, while traffic activity is updated on per A-MPDU basis.

Handoff trigger: MAPS associates a timer with each client $k \in K$. It triggers AP selection for k, when its timer expires, or when a special event occurs. Since handoff process lasts approximately for 1.5 seconds in our testbed, MAPS sets client's timer in the order of tens of seconds. It freezes the timer for very low traffic clients (i.e., $\overline{S}_{ampdu} \approx 0$), to prevent frequent handoffs. It also defers handoff process, when delay-sensitive traffic (e.g., VoIP) is in progress. Client assignment is also triggered upon client's mobility. MAPS detects mobility through SNR variations. By using both timers and events for handoff trigger, MAPS remains adaptive to channel dynamics, without triggering excessive handoffs.

Group formation: Upon triggering handoff for client k, MAPS will identify the set of APs in its range. To reduce the MU-MIMO group formation computational overheads, MAPS excludes in advance APs with very low RSSI, since they will likely be sub-optimal in terms of throughput. Then, it calculates the set of groups G'_α (and $G'_{k,\alpha}$) for APs in k's range, using the greedy search approach proposed in [21]. In summary, it first sorts in descending order, the clients based on their SU-MIMO throughput, and iteratively goes through the list, to group the clients that provide the highest aggregate throughput with the already selected clients. The search terminates when the group is complete, or when adding more clients to a group results in lower throughput than serving them in SU-MIMO mode. Finally, MAPS computes k's throughput at AP α from equation 5. It assigns k to the AP which maximizes its throughput.

Our experimental results show that MAPS heuristic algorithm performs very close to the optimal AP selection.

4.4 Discussion

Downlink throughput: MAPS seeks to improve downlink network throughput. This is because 802.11ac MU-MIMO is only supported in downlink direction, which dominates the uplink, in Wi-Fi networks [18]. However, improving downlink throughput, allows for more airtime (and hence better throughput) in uplink transmissions as well.

Co-existence with existing AP functions: MAPS can work in concert with existing AP core functionalities, such as MU-MIMO grouping and rate adaptation algorithms. MAPS' SINR estimation could be also used for even improving such algorithms. We leave such extensions for future work.

Legacy clients: MAPS seeks to optimize the assignment of IEEE 802.11ac MU-MIMO-capable clients to APs. Legacy 802.11a/b/g/n clients are assigned to APs, based on legacy algorithms. However, MAPS' throughput model still considers the load offered by legacy clients connected to an AP.

5 IMPLEMENTATION

Our AP's firmware has access to only 960 KB on chip memory, whose 98% is already occupied by various functions. Given that one decompressed CSI sample requires 3.6 KB memory space, storing multiple CSIs for all connected clients in AP's firmware is not feasible. Constructing CSI profiles in the more powerful AP's "host" board, which runs on a dual core 1.4 GHz CPU with a 512 MB DDR3 memory may overload the AP. Instead, MAPS constructs clients' CSI profiles at the controller, considering that all data packets are going through the controller and CSI communication overhead is small. However, due to their low memory requirements, we still implement clients' bandwidth and traffic profiles, and maintain channel busy time statistics, in AP's firmware. The rest of MAPS' functionality is implemented (using MATLAB) at the controller, as shown in Figure 4. Our controller is a laptop with 8GB DDR3 memory and 4 Intel core i7-3520M CPU. Note that our implementation does not include the MLME module and hence does not support real time handoff. To this end, upon estimating the best AP-client assignments, we manually assign clients to these APs, and then evaluate the network's performance.

6 EVALUATION

In this section, we evaluate MAPS, using testbed experiments and trace-driven simulations. We compare MAPS with DenseAP [17], which is representative of legacy AP selection designs proposed by research studies, and deployed by AP vendors. DenseAP uses an available capacity metric, which estimates the throughput as a function of PHY rate (calculated from RSSI) and Wi-Fi channel

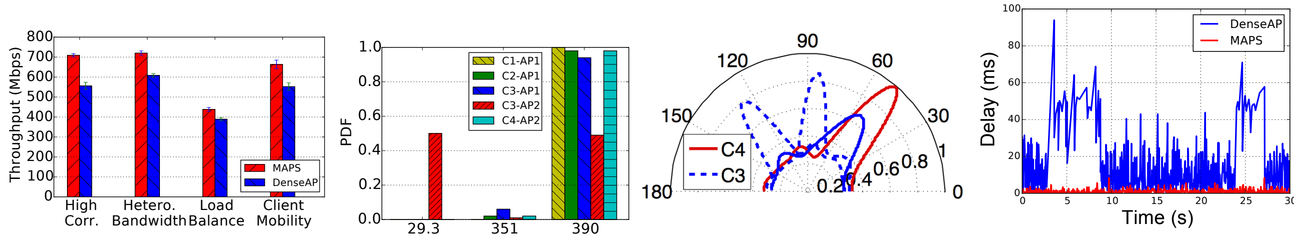

(a) Aggregated UDP throughput for MAPS and DenseAP.
(b) Clients' rate distribution in MU-MIMO.
(c) Dominant multipaths for mobile C3 and static C4, at AP2.
(d) MAPS vs. DenseAP over VoIP.

Figure 8: MAPS and DenseAP performance in representative settings. MAPS performs similar to Oracle.

busy time. It seeks to balance the load by assigning clients to less loaded APs[5]. We also compare MAPS with an "Oracle", which is the best-throughput (optimal) client assignment. Oracle finds the best setting through exhaustive search. Our experimental setup consists of the 802.11ac APs and phones described in Section 3.1. We evaluate multiple topologies, under various traffic scenarios (UDP, TCP, VoIP).

6.1 Performance in Representative Settings

We first evaluate MAPS' performance in four representative settings, which capture different aspects of dynamics in 802.11ac networks. We consider that APs generate saturated downlink UDP traffic to clients, and that they operate on orthogonal channels, unless it is explicitely mentioned.

Correlated channels: We first evaluate MAPS in our case study setting of Figure 3. Similar to Oracle, MAPS can identify the best-throughput client assignment, and achieves 153 Mbps (or 27.5%) aggregated throughput gain compared to DenseAP, as shown in Figure 8a. Specifically, MAPS can identify the high correlation of C3, C4's wireless channels at AP2, by computing their CSI correlation and SINR values. The correlation factor $\rho(C3, C4)$ is 41% and 50% higher compared to $\rho(C1, C3)$ and $\rho(C2, C3)$, as shown in Table 3. Such high correlation does not allow for MU-MIMO operation at AP2. Hence, MAPS will assign C3 to AP1.

DenseAP will falsely assign C3 to the highest RSSI AP2. Although the SU-MIMO is the best mode for such assignment (cf. Tab. 1), we observed that the AP's MU-MIMO client grouping algorithm will periodically try to evaluate the performance of the group $\{C3, C4\}$. However, this will result in high inter-client interference and high PER (cf. Fig. 2b). When C3, C4 are grouped together, AP's PHY rate adaptation switches to low PHY rates to cope with such interference. This is shown in the rate distribution Figure 8b, where the AP2 often uses the lowest available 802.11ac rate (29.3 Mbps) to transmit to C3. However, when MAPS assigns C3 to the lower RSSI AP1, the selected PHY rate is mostly 390 Mbps, which results in higher throughputs.

Table 3: CSI correlation for case study setting.

Setting	C1-C2	C1-C3	C2-C3	C3-C4
CSI Corr	0.42	0.51	0.48	0.72

Heterogeneous bandwidths: We next evaluate MAPS with heterogeneous bandwidth clients. We deploy two APs and five clients. C1 and C2 are connected to AP1 operating at 80 MHz. C4 and C5 are connected to AP2. Due to interferences, C4 and C5 operate at 40 MHz. A new client C3 has a stronger RSSI with AP2, than AP1. Hence, DenseAP will assign C3 to AP2 without considering that, C3 cannot form an 80 MHz MU-MIMO group at AP2. C3's throughput is 79 Mbps at AP2, while the total network throughput is 608 Mbps (cf. Fig. 8a). However, MAPS can identify the opportunity of a high throughput MU-MIMO group {C1, C2, C3} at 80 MHz, and assigns C3 to AP1. Associating with AP1, C3 achieves 156 Mbps throughput, with a total network throughput of 720 Mbps. Leveraging client's bandwidth profile, MAPS can almost double C3's throughput. It also boosts total network throughput by 112 Mbps (18.4%) compared to DenseAP. MAPS performs the same as Oracle.

Mobility: We next study MAPS' responsiveness to mobility. We deploy two APs, three static and one mobile client. Static clients C1, C2 are associated with AP1, and C4 with AP2. C3 is moving with pedestrian speed around AP2. Our traces show that the highest RSSI AP for C3 is always AP2. Hence, DenseAP assigns C3 to AP2. This results in 65 Mbps throughput for C3 and 552 Mbps network throughput.

MAPS monitors the channel dynamics of the mobile client C3, and constructs a CSI profile for AP2 with three dominant multipaths, as shown in Figure 8c. The most dominant path among the three is represented with the solid line. Interestingly, C3 and C4 channels overlap in space at AP2 (cf. Fig. 8c), which implies correlated channels and high inter-client interference. To avoid client groups with correlated channels, MAPS will assign C3 to AP1. C3 achieves 91 Mbps throughput at AP1, and the aggregated network throughput is 663 Mbps (cf. Fig. 8a). This corresponds to a throughput gain of 40% for C3, compared to DenseAP. Network throughput is also increased by 20.1%.

Unbalanced traffic: Different from the previous experiments, we next evaluate a setting where two APs operating on the same channel, have unbalanced loads. Specifically, clients C1 and C2 are connected to AP1, while AP2 does not serve any client. Let's now consider a new client C3 in the network, with similar RSSI from AP1 and AP2. DenseAP will assign C3 to AP2, to balance the load across APs. Such assignment results in 75 Mbps and 389 Mbps throughput for C3 and for the network, respectively. On the other hand, MAPS will estimate factor $w_{C3,AP1}$ to be equal to 1, and $w_{C3,AP2}$ to be 1/2. Hence, given a negligible inter-client interference among C1, C2, C3 at AP1, MAPS will assign C3 to AP1. Such assignment almost doubles C3's throughput. It also increases the network throughput by 50 Mbps.

Delay-sensitive traffic: We finally evaluate MAPS over delay-sensitive traffic, such as VoIP. In our setting, clients C1, C2 are connected to AP1, and C4, C5 to AP2. Both APs generate saturated downlink UDP traffic to clients. DenseAP assigns a new client C3 to the highest-RSSI AP1, while MAPS connects C3 to AP2, which maximizes MU-MIMO gains. Then, C3 initiates a VoIP call to another device connected to the AP through Ethernet. Figure 8d shows the one-way network delay for VoIP traffic over a 30-second time window, for C3 at AP1 (DenseAP), and C3 at AP2 (MAPS). We observe

[5]DenseAP performs transmit power control, which is out of the scope of this work.

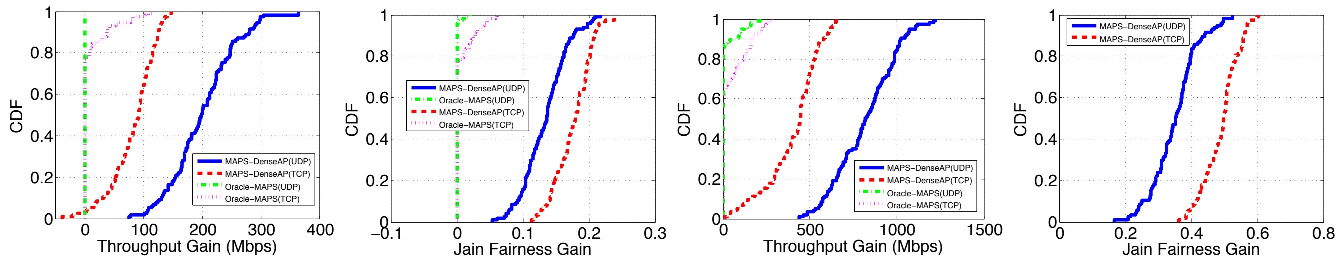

(a) Aggregated throughput difference for field trials.

(b) Jain fairness index difference for field trials.

(c) Aggregated throughput difference for simulations.

(d) Jain fairness index differencen for simulations.

Figure 9: Throughput and fairness comparison of MAPS, Oracle and DenseAP.

that the average and peak delays are only 1.1 ms and 8.4 ms, when C3 connects to AP2. This is because MAPS' assignment allows for C3, C4, C5 to form an low interference MU-MIMO group. However, DenseAP assignment results in 11.4× higher average delay (12.4 ms) compared to MAPS. For 7% of the samples, the delay for DenseAP exceeds 30ms, which is above the delay requirements of VoIP applications [16]. This is because C3 mainly operates in SU-MIMO mode at AP1, due to high inter-client interference with C1 and C2. Interestingly, Figure 8d shows high delay variations for DenseAP assignment, with delay peaks up to 94ms. We observe that such peaks (at [2.7, 8.8] sec. and [23.8, 27.2] sec.) appear, when the MU-MIMO grouping algorithm tries to group C1, C2, C3 together. Such grouping creates high PER and low throughput.

6.2 Larger Scale Field Trials

We further experimentally evaluate MAPS in multiple larger scale topologies with 6 APs and 20 clients. We present the experimental floorplan in our technical report [28]. Apart from our APs, we detect 22 more BSSIDs at 5 GHz, to operate in various channels, in the same RF coverage zone. We run each experiment for 5 minutes at different times of day considering both static and mobile clients, and we report experiments from multiple runs. For each run, we compare MAPS, DenseAP and Oracle, for saturated downlink UDP and single-flow TCP traffic. Figure 9a shows the aggregated (over all APs and clients) network throughput gains of MAPS over DenseAP. Each point of the distribution reflects a different setting. We observe that MAPS performs similar or better than DenseAP in more than 90% of the settings. For UDP, the gain is at least 149 Mbps in 50% of the settings, and it can go up to 365 Mbps (which corresponds to a 52.3% gain). Throughput gains for TCP are smaller (23.6%). This is because UDP traffic is always saturated compared to TCP. The highest gains for MAPS are observed in static client settings, when clients' channels are highly correlated. The smallest gains (or even loss) for MAPS are observed in highly dynamic environments, where MAPS' CSI profile may not capture the channel dynamics, and may assign clients to lower RSSI APs, which happen to be the lower throughput APs. In such settings, DenseAP outperforms MAPS by up to 95 Mbps (cf. Fig. 9a).

MAPS mostly performs similar to Oracle. Specifically, Figure 9a shows that MAPS throughput is the same with Oracle in 94% of the settings for UDP, and 75% of the settings for TCP. Oracle outperforms MAPS in a few scenarios of highly dynamic environments, as we discussed above.

Interestingly, MAPS can also improve the throughput fairness among clients connected to APs in the same vicinity. We illustrate our finding in Figure 9b, which plots the difference of Jain Fairness Index between MAPS and DenseAP, and between Oracle and MAPS, for UDP and TCP. MAPS has always equal or larger Jain index compared to DenseAP. The difference exceeds 0.2, which is a significant gain, if we consider that an index of 1 implies perfect fairness. Since MAPS limits inter-client interferences (and hence PER), it does not negatively affects certain clients' TCP windows. This results in better TCP fairness gains (compared to UDP) over DenseAP (cf. Fig. 9b). MAPS achieves the same fairness as Oracle in the vast majority of the settings.

6.3 Trace Driven Simulations

We next conduct trace-driven simulations to evaluate MAPS in larger scale network topologies. For our simulation, we have collected wireless link performance traces (e.g., CSI, throughput, PER) from multiple settings. We have also collected per-client traffic load statistics from 82 APs of an enterprise Wi-Fi network, to simulate realistic traffic scenarios. We then combine these traces to simulate larger networks. We simulate scenarios where 20 APs and 108 (static and mobile) clients are placed in a building floor. The number of clients per AP varies from 1 to 54.

In Figure 9c, we present the aggregated (over all APs and clients) network throughput gains of MAPS over DenseAP. We observe that MAPS always performs similar or better than DenseAP. Specifically, it achieves up to 1.2 Gbps (or 28.6%) and 663.8 Mbps (or 17.5%) throughput gain over UDP and TCP, respectively. Figure 9c, further shows that MAPS performs the same as Oracle, for 85% and 60% of the settings, for UDP and TCP, respectively.

Interestingly, our results show that MAPS' gains over DenseAP do not necessarily drop when the number of clients connected to an AP increases. For example, Table 4 shows the average throughput for MAPS, DenseAP and Oracle, when the number of clients in two APs' vicinity varies from 12 to 108. We observe MAPS' gains over DenseAP range from 33% to 45%, with the maximum gain to be achieved for 108 clients. This is because, the higher number of clients connected to an AP does not necessarily increase the MU-MIMO grouping opportunities. Specifically, we observe that the number of candidate clients for grouping at each transmit opportunity (TXOP) is limited by: a) the active clients, b) the fair scheduler implemented in our AP, which will not reschedule the clients served in the previous TXOPs, c) the clients' correlated channels and channel bandwidth configurations. Particularly, we observe that the MU-MIMO groups' size for DenseAP is typically less than maximum supported size of 3 clients, or it often operates in SU-MIMO.

Finally, our simulations verify that MAPS can improve the fairness among clients, as shown in Figure 9d. In conclusion, our experiments show that MAPS can significantly boost the performance for large Wi-Fi networks.

Table 4: Average throughput for varying number of clients.

Total clients	12	36	60	84	108
MAPS Thr. (Mbps)	313 ± 5	314 ± 2	333 ± 1	341 ± 1	346 ± 1
MAPS clients AP1/AP2	6/6	19/17	31/29	40/44	52/56
DenseAP Thr. (Mbps)	223 ± 5	223 ± 2	250 ± 1	242 ± 2	237 ± 1
DenseAP clients AP1/AP2	6/6	18/18	32/32	42/42	54/54
Oracle Thr. (Mbps)	360 ± 3	361 ± 2	360 ± 2	362 ± 1	362 ± 1
Oracle clients AP1/AP2	6/6	19/17	31/29	40/44	52/56

7 RELATED WORK

There are several studies related to our work.

AP selection: AP selection algorithms can be classified in centralized [17, 29] and distributed [11, 13, 25]. Similar to MAPS, in centralized solutions, APs exchange RSSI, traffic load, interference feedback with a controller, which decides the network-wide optimal client assignment. In distributed algorithms, it is the client which selects the best AP. However, all the above systems have been designed for legacy 802.11a/b/g/n networks and are oblivious to MU-MIMO feature. Hence, they can limit the MU-MIMO grouping opportunities, as shown by our experiments. AP selection designs proposed by AP vendors [2, 8], are also RSSI-based and have the same limitations in MU-MIMO settings.

The theoretical study in [9] jointly solves the problems of MU-MIMO AP selection and client grouping. It seeks to assign clients with uncorrelated channels to the same AP. However, such proposal has two key limitations. First, it is oblivious to clients' bandwidth configurations, and it does not consider the impact of AP load and channel utilization to throughput performance. Hence, it performs poorly in the scenarios described in Sections 3.3, 3.4. Second, our results have shown that MU-MIMO grouping and AP selection happen at different time scales (msec. and sec. scales, respectively). Thus, triggering AP selection at msec. granularity can cause excessive handoff overheads. Different from [9], MAPS decouples these two functions, and considers clients' heterogeneity and AP load, when assigning clients to APs.

MU-MIMO grouping and scheduling: There have been several MU-MIMO client grouping and scheduling proposals [19, 21, 23, 30]. Such designs can only achieve high MU-MIMO gains, if clients with uncorrelated channels have been assigned to an AP. Thus, they can realize their full potential, only by working in concert with designs like MAPS. Moreover, MU-MIMO grouping designs leverage explicit beamforming feedback to identify uncorrelated channels. Such approach requires excessive clients' handoffs (cf. Sec. 4.1), and it is not efficient for AP selection. Hence, MAPS uses implicit CSI feedback to assign clients to APs.

Network MU-MIMO: MAPS assigns clients with orthogonal channels to APs, to allow for MU-MIMO groups with no inter-client interference. For a given client assignment, recent designs [4, 14, 24, 26] enable APs and clients in interfering cells to coordinately cancel the inter-cell interference, using their antennas for beamforming and interference cancellation. Such solutions typically require client-side modifications and are not 802.11-compliant. On the other hand, we implement MAPS in 802.11ac-compliant commodity APs and controllers. Note that, MAPS could also work in concert with network MU-MIMO, to further improve performance.

8 CONCLUSION

In this paper, we have studied the AP selection problem in MU-MIMO Wi-Fi networks, using commodity 802.11ac testbeds. Our experimental results show that legacy AP selection designs assign clients with correlated channels and heterogeneous bandwidths to the same AP, limiting the MU-MIMO grouping opportunities. Their approach to load balancing is also MU-MIMO oblivious and can decrease the MU-MIMO gains. To this end, we propose a new Mu-mimo-Aware AP Selection (MAPS) design, which can identify the best-throughput client assignment, at low overhead. Our results show that MAPS significantly outperforms legacy designs. We believe that MAPS can be a key building block for designing the future MU-MIMO 802.11ax and 5G networks.

REFERENCES

[1] Virtual cells: The only scalable multi-channel deployment. In *Meru Networks White Paper*, 2005.
[2] Clientmatch technology: Wi-fi client optimization. In *Aruba Networks White Paper*, 2015.
[3] Ericsson 5G field trial gear achieves peak downlink throughput over 25 Gbps with MU-MIMO. In *Ericsson press release*, 2016.
[4] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire. Achieving high data rates in a distributed mimo system. In *ACM MobiCom'12*.
[5] W. S. Cleveland. Robust locally weighted regression and smoothing scatterplots. In *Journal of the ASA*, volume 74, pages 829–836, 1979.
[6] M. X. Gong, B. Hart, and S. Mao. Advanced Wireless LAN Technologies: IEEE 802.11ac and Beyond. In *ACM GetMobile: Mobile Comp. and Comm.*, 2015.
[7] D. Halperin, W. Hu, A. Sheth, and D. Wetherall. Predictable 802.11 packet delivery from wireless channel measurements. In *ACM SIGCOMM'10*.
[8] S.-w. Han and M. Ram. Distributed client steering algorithm to a best-serving access point. In *USA Patent, US9078198, Meru Networks*, 2015.
[9] Y.-C. Hsu, K. C.-J. Lin, and W.-T. Chen. Client-ap association for multiuser mimo networks. In *IEEE ICC'15*.
[10] IEEE Standards Association. IEEE Standards 802.11ac-2013: Enhancements for Very High Throughput for Operation in Bands below 6 GHz. 2013.
[11] T. Iwami, Y. Takaki, K. Yamori, C. Ohta, and H. Tamaki. Distributed association control considering user utility and user guidance in ieee 802. 11 networks. In *IEEE PIMRC'13*.
[12] R. Jain, D.-M. Chiu, and W. R. Hawe. *A quantitative measure of fairness and discrimination for resource allocation in shared computer system*. DEC TR, 1984.
[13] H. Lee, S. Kim, O. Lee, S. Choi, and S.-J. Lee. Available bandwidth-based association in ieee 802.11 wireless lans. In *ACM MSWiM'08*.
[14] K. C.-J. Lin, S. Gollakota, and D. Katabi. Random access heterogeneous mimo networks. In *ACM SIGCOMM'11*.
[15] C. Lukaszewski. Very high-density 802.11ac networks, planning guide. In *White Paper, Aruba Networks*, 2014.
[16] Microsoft. Media quality and network connectivity performance in skype for business online. In *https://support.office.com*.
[17] R. Murty, J. Padhye, R. Chandra, A. Wolman, and B. Zill. Designing high performance enterprise wi-fi networks. In *USENIX NSDI'08*.
[18] I. Pefkianakis and et al. Characterizing home wireless performance: The gateway view. In *IEEE INFOCOM'15*.
[19] W.-L. Shen, K. C.-J. Lin, M.-S. Chen, and T. Kun. SIEVE: Scalable User Grouping for Large MU-MIMO Systems. In *IEEE INFOCOM'15*.
[20] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong. Argos: Practical many-antenna base stations. In *ACM MobiCom'12*.
[21] S. Sur, I. Pefkianakis, Z. Xinyu, and K.-H. Kim. Practical mu-mimo user selection on 802.11ac commodity networks. In *ACM MobiCom'16*.
[22] W. Wang, A. X. Liu, M. Shahzad, K. Ling, and S. Lu. Understanding and modeling of wifi signal based human activity recognition. In *ACM MobiCom'15*.
[23] X. Xie and X. Zhang. Scalable user selection for MU-MIMO networks. In *IEEE INFOCOM'14*.
[24] X. Xie, X. Zhang, and E. Chai. Cross-cell dof distribution: Combating channel hardening effect in multi-cell mu-mimo networks. In *ACM MobiHoc'15*.
[25] S. Xu, F. Ren, Y. Xu, C. Lin, and M. Yao. Selecting a preferable access point with more available bandwidth. In *IEEE ICC'13*.
[26] H. Yu, O. Bejarano, and L. Zhong. Combating inter-cell interference in 802.11ac-based multi-user mimo networks. In *ACM MobiCom'14*.
[27] Y. Zeng, P. H. Pathak, and P. Mohapatra. A first look at 802.11 ac in action: Energy efficiency and interference characterization. In *IFIP Networking'14*.
[28] Y. Zeng, I. Pefkianakis, K.-H. Kim, and P. Mohapatra. Maps: Mu-mimo-aware ap selection for 802.11ac networks. In *http://spirit.cs.ucdavis.edu/pubs/tr/MAPS-TR.pdf*.
[29] Y. Zhu, Q. Ma, C. Bisdikian, and C. Ying. A user-centric network management framework for high-density wireless lans. In *IFIP/IEEE IM'09*.
[30] G. Zois, A. Michaloliakos, K. Psounis, V. Vassalos, and I. Mourtos. Non-asymptotic performance bounds for downlink mu-mimo scheduling. In *IEEE/IFIP WONS'16*.