

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

12-2014

MIDAS: Empowering 802.11ac with multiple-input distributed antenna systems

Jie XIONG

Singapore Management University, jxiong@smu.edu.sg

Karthikeyan SUNDARESAN

University College London

Kyle JAMIESON

University College London

Mohammad A. KHOJASTEPOUR

NEC Laboratories America

Sampath RANGARAJAN

NEC Laboratories America

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Software Engineering Commons](#)

Citation

XIONG, Jie; SUNDARESAN, Karthikeyan; JAMIESON, Kyle; KHOJASTEPOUR, Mohammad A.; and RANGARAJAN, Sampath. MIDAS: Empowering 802.11ac with multiple-input distributed antenna systems. (2014). *CoNEXT'14: Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies: December 2-5, 2014, Sydney, Australia*. 29-40.

Available at: https://ink.library.smu.edu.sg/sis_research/2699

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

MIDAS: Empowering 802.11ac Networks with Multiple-Input Distributed Antenna Systems

Jie Xiong[†], Karthikeyan Sundaresan[‡], Kyle Jamieson[†], Mohammad A. Khojastepour[‡],
Sampath Rangarajan[†]

[†]Department of Computer Science, University College London, UK

[‡]NEC Laboratories America, NJ, USA

{j.xiong, k.jamieson}@cs.ucl.ac.uk[†], {karthiks, amir, sampath}@nec-labs.com[‡]

ABSTRACT

Next generation WLANs (802.11ac) are undergoing a major shift in their communication paradigm with the introduction of multi-user MIMO (MU-MIMO), transitioning from single-user to multi-user communications. We argue that the conventional AP deployment model of co-located antennas as well as their PHY and MAC mechanisms are not designed to realize the complete potential of MU-MIMO. We propose to leverage distributed antenna systems (DAS) to empower next generation 802.11ac networks. We highlight the multitude of benefits that DAS brings to MU-MIMO and 802.11ac in general. However, several challenges arise in the process of realizing these benefits in practice, where avoiding client modifications and making only minimal software modifications to APs is important to enable rapid adoption. Towards addressing these challenges, we present the design and implementation of *MIDAS*, the *Multiple-Input Distributed Antenna System*. MIDAS couples a DAS deployment of AP antennas with a suite of novel yet standards-compatible mechanisms at the PHY and MAC layers that best leverage the DAS deployment to maximize 802.11ac performance. Our WARP-based experimental evaluation demonstrates MIDAS's ability to significantly boost the performance of current 802.11ac design, demonstrating throughput gains over 802.11ac MU-MIMO for 100-200%, while remaining amenable to commercial adoption.

Categories and Subject Descriptors

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

Keywords

MIDAS; Distributed antenna system; DAS; Multi-user MIMO; MU-MIMO; 802.11ac; Multiple input

1. INTRODUCTION

Recently, wireless LAN (WLAN) designs are undergoing a paradigm shift, evolving from single-client to multi-client communication patterns with the introduction of the 802.11ac standard. One of the key reasons behind this move is a form factor of smart devices that limits


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

CoNEXT'14, December 2–5, 2014, Sydney, Australia.

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

ACM 978-1-4503-3279-8/14/12 ...\$15.00.

<http://dx.doi.org/10.1145/2674005.2675014>



(a) A Co-located Antenna System (CAS)

(b) A Distributed Antenna System (DAS)

Figure 1: CAS vs DAS deployment.

the number of antennas to one or two, with a consequent limit on the number of streams that the AP can send to a single client, even if the AP has many more antennas available. Multi-user MIMO allows the AP to use all its antennas to send multiple streams of data to different clients simultaneously. Clients measure the channel to the AP and feed this *channel state information* (CSI) back to the AP, which in turn sends different data streams on the downlink such that each client receives only the data stream it needs, while other unwanted (and otherwise interfering) data streams are suppressed (this process is called *precoding*). A popular and lightweight precoding scheme is called *zero-forcing beamforming* [27] (ZF-BF), wherein undesired data streams are nulled (forced to zero) at a client. MU-MIMO thereby allows the AP to better leverage its spatial (antenna) resources and achieve multiplicative throughput increases: this is called *multiplexing gain*.

While the incorporation of MU-MIMO in 802.11ac is a welcome step towards utilizing the AP's radios and antennas effectively, the current WLAN deployment model is to co-locate all an AP's antennas within a few wavelengths of the carrier frequency, as shown in Figure 1(a); we refer to this as a co-located antenna system, or CAS. In this work we show that spatially separating the antennas of an 802.11ac AP via RF or optical cabling, and leveraging MU-MIMO through a *Distributed Antenna System* (DAS), as shown in Figure 1(b), is essential to getting the most out of MU-MIMO.

Specifically, a DAS design has the potential to improve WLAN performance by means of three distinct mechanisms:

1. Spatial Diversity. As illustrated in Figure 1, distributing the AP's antennas helps a client find an AP antenna closer to it on average compared to in a CAS. This reduces the path loss on average and hence improves link SNR and capacity.

2. Cell Capacity. In MU-MIMO, when data streams are sent to different clients simultaneously from the AP, they interfere with each

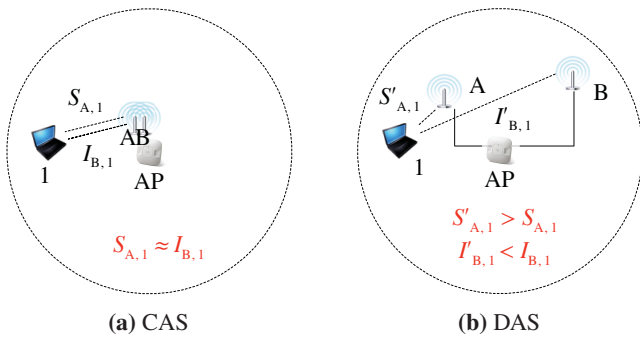


Figure 2: DAS increases the signal strength but also decreases the interference cancellation burden on precoding.

other. This mutual interference is addressed through data encoding (precoding) at the AP. With DAS, a client will likely have a much stronger signal ($S'_{A,1}$) and weaker interference ($I'_{B,1}$) as shown in Figure 2(b). With carefully-chosen precoding, this better isolates desired and interfering streams compared to a CAS system, where path loss from different antennas is almost the same, as illustrated in Figure 2(a). The result is an improved effective SINR and hence better network throughput.

3. Network Capacity. In a DAS, the channel states (busy or idle) at different AP antennas can differ, unlike in an 802.11ac CAS. With each of the antennas having the ability to sense and decouple the channel states helps leverage spatial reuse with finer granularity. This not only allows for packing more concurrent transmissions in the network, resulting in improved network capacity, but also alleviates the hidden terminal issue (exacerbated with multi-client communication) and reduces the deadzones in the coverage area.

Challenges in realizing DAS. While DAS-based MU-MIMO systems might appear similar in effect to deploying multiple CAS APs (in place of distributed antennas), there are several performance and practical benefits to the former that directly cater to 802.11ac as we discuss in Section 2. To realize DAS, three practical, yet technically challenging issues arise:

Firstly, while a DAS’s distributed antennas can emphasize the desired streams at closer clients more than the interfering streams, a conventional MU-MIMO precoding scheme such as zero-forcing beamforming (ZFBF) that splits the power equally across the streams neither meets practical 802.11ac per-antenna power constraints in a DAS nor leverages the inherent topology advantage of a DAS. On the other hand, precoding schemes for distributed MU-MIMO systems with a per-antenna power constraint are either too computationally complex to realize [11, 32] or require modifications to client receivers, hindering adoption.

Secondly, since it is designed for a CAS, 802.11ac’s MAC keeps a single medium access channel state (idle or busy) for the entire AP, but a DAS’ antennas are distributed, hence have different true channel states. This channel state coupling significantly limits the AP’s ability to achieve fine-grained spatial reuse and alleviate hidden terminals in the network.

Finally, the MAC in a CAS treats all antennas alike, employing all during the downlink precoded transmission. But unlike in a CAS, in a DAS there is a large disparity across the channel gains from each antenna to a client. Hence, selecting clients intelligently based on the set of available idle antennas for MU-MIMO has a much larger impact in DAS and has to be carefully tailored.

MIDAS. In this paper we propose MIDAS (the *Multiple Input Distributed Antenna System*), a WLAN design that combines a DAS deployment of antennas with several software mechanisms at the

PHY and MAC layer to leverage the full potential of MU-MIMO. MIDAS is standards-compliant in that it does not require any modifications to the PHY, requires minimal software modifications to the AP’s MAC, while maintaining complete transparency to clients.

At the PHY layer, MIDAS’s design includes a novel, lightweight, *power-balancing* mechanism for ZFBF-based precoding. It respects 802.11’s per-antenna power constraint, while leveraging the inherent topology imbalance between desired and interfering data streams to appropriately assign power to the different streams. This enables MIDAS to deliver near-optimal precoding performance. Its lightweight nature is conducive for implementation especially for channels with small channel coherence time, and transparency to clients fosters easier deployment and adoption. Further, with topology imbalance being an inherent challenge for distributed MU-MIMO systems in general, our precoding scheme also improves upon those in the literature [19, 32].

At the MAC layer, MIDAS incorporates two novel mechanisms. First, it re-designs the carrier sensing mechanism to allow for *fine-grained carrier sensing* at each of the AP’s antennas by maintaining multiple, independent medium occupancy states. With fine-grained carrier sensing enabled, MIDAS allows for improved spatial reuse and also alleviates hidden terminals. Second, it enables *virtual packet tagging*, where each packet (of a client) is tagged with a preferred list of antennas for transmission. The latter provides a means to perform client filtering for MU-MIMO transmissions that is tailored to the set of antennas available. This enhances MU-MIMO performance considerably, while allowing the MAC to leverage fine-grained spatial reuse opportunities.

We have implemented MIDAS on the Rice WARP [20] platform, and deployed DAS with RF cables in two different indoor office testbed environments. Experimental results demonstrate capacity gains of around 30% from MIDAS’s precoding, 40% from DAS deployment and 65% from its MAC mechanisms, to result in a net gain of up to 200% over conventional 802.11ac MU-MIMO systems. These results show that it is possible to significantly boost the performance of MU-MIMO in next-generation WLANs while being amenable to practical realization with off-the-shelf clients.

The rest of this paper is structured as follows. Following a quick background on the most relevant technologies (§2), we present the design of MIDAS in §3, highlighting the above novel contributions and explaining how MIDAS addresses the above challenges. §4 describes our implementation of MIDAS. An experimental evaluation (§5) follows, investigating how overcoming the above challenges improves performance and why. §6 surveys the related work. §7 presents some additional discussion, and §8 concludes.

2. BACKGROUND

Conventional DAS. The distributed antennas of a DAS system merely radiate the received RF signal (from the AP) within the AP’s contention domain. WLAN [7] and mobile cellular networks [4, 29] have long used conventional DASs to provide improved wireless coverage indoors, especially in larger venues, such as convention centers or stadiums. A conventional DAS simply broadcasts the *same* signal from all antennas with some designs even using “leaky feeder” coaxial cable that radiates RF energy continuously along the length of the cable [21]. Thus conventional DASs preclude the use of MU-MIMO and its associated spatial multiplexing gains. In contrast, MIDAS is a true MU-MIMO DAS that can send different RF signals to different antennas (i.e multiple input).

Multi-AP distributed MU-MIMO. Here transmitting antennas are located at different APs, which run a fine-grained synchronization protocol for phase-coherent transmissions. Considering the dif-

ference between MIDAS and a multi-AP distributed MU-MIMO system, there are several key differences. Distributed MIMO systems entail phase and time synchronizing all the APs and sharing all their data on the wired backhaul network to enable distributed MU-MIMO (e.g. MegaMIMO [19], AirSync [5]), which is a practical impediment to realization and deployment especially if different organizations or people administer different APs. Even if feasible, it would require a revamp of the 802.11ac protocol to allow multiple 802.11ac APs to share and jointly precode data to multiple clients. In MIDAS, each 802.11ac AP serves as the central processing node to realize MU-MIMO across its distributed antennas. Hence, phase and time synchronization comes for free, and incurs no backhaul overhead (RF signals are sent over DAS cables).

Deploy multiple conventional 802.11 CAS APs. Consider a Multi-AP approach, where a CAS AP is deployed in place of each of the antennas of a MIDAS deployment. Being logically similar, there are several key differences in both practical realization and performance between the two. Comparing MIDAS with Multi-AP, where each of the CAS APs is a *multiple* antenna AP, is unfair to MIDAS in terms of antenna and AP resources. However, even in this case, MIDAS has an advantage over Multi-AP. Note that the antennas of a MIDAS AP are distributed within the contention domain of the AP. Deploying multiple APs within one content domain is inefficient. The counterpart CAS APs in Multi-AP are still restricted to MU-MIMO transmissions from a single CAS AP since 802.11ac does not allow for cooperation across APs. On the other hand, an multiple stream transmission of higher aggregate rate is possible in MIDAS due to the potentially higher rank channel matrix (due to un-correlated channels) created by the distributed antennas. Further, when each CAS AP consists of a single antenna in Multi-AP (for a fair comparison), it would be restricted to a single stream within a contention domain unlike MIDAS that can support multiple streams.

In essence, it is better to distribute the antenna resources of an AP within a single contention domain for multi-user transmissions, and DAS provides a practical primitive to realize this. DAS confers the benefits of distributed MU-MIMO to a CAS-based MU-MIMO system, while remaining compatible with 802.11ac, thereby fostering easier adoption. However, it requires appropriate precoding and MAC mechanisms to deliver its gains, which in turn forms the objective for MIDAS.

3. DESIGN

In this section we describe the MIDAS design elements that complement the DAS model with a suite of PHY- and MAC-layer mechanisms capable of delivering the most benefits out of a multi-user communication system, while remaining standards-compliant. We begin by considering how to precode transmissions at the AP's distributed antennas while managing transmission power (§3.1). Next we discuss MIDAS's fine-grained channel sensing MAC (§3.2.2) and client selection algorithms (§3.2.4, §3.2.5) that are tailored for the DAS deployment. We put everything together in §3.2.1 and conclude with a discussion of how the above design elements fit readily into 802.11ac (§3.3), fostering easier adoption.

3.1 Pre-coding for MU-MIMO in MIDAS

To place MIDAS's precoding algorithm into context, we begin with a brief primer on conventional ZFBF, explaining the per-antenna power constraint that 802.11ac places on each of the AP's antennas and the reasons behind this constraint.

3.1.1 Primer: Zero-forcing beamforming

ZFBF maximizes the SNR of each desired stream subject to the constraint that the interference each causes to other clients be completely

nulled. Such an approach is optimal for a total power constraint [28] and can be formulated as an optimization problem to choose precoding vectors $\vec{v}_j = \{v_{1j} \cdots v_{|\mathcal{T}|j}\}^T$ for every client j in the set of chosen clients \mathcal{C} , where \mathcal{T} is the set of transmit antennas.¹ The objective is to maximize the sum rate to the set of chosen clients \mathcal{C} :

$$\arg \max_{\vec{v}_j} \sum_{j=1}^{|\mathcal{C}|} \log(1 + \rho_j) \quad (1)$$

$$\text{where } \rho_j = \frac{|\sum_{k=1}^{|\mathcal{T}|} h_{jk} v_{kj}|^2}{N_o} \quad (\forall j \in \mathcal{C})$$

subject to the following two constraints:

$$\sum_{k=1}^{|\mathcal{T}|} \sum_{j=1}^{|\mathcal{C}|} |v_{kj}|^2 \leq P_{tot} \quad (\forall k \in \mathcal{T}) \quad (2a)$$

$$\sum_{k=1}^{|\mathcal{T}|} h_{ik} v_{kj} = 0 \quad (j \in \mathcal{C}, i \neq j) \quad (2b)$$

where h_{jk} is the complex channel gain from antenna k to client j . Equation (2a) is the total power constraint, and (2b) is the precoding constraint, i.e., precoded symbols should cancel each other at client i while arriving at a desired client j ($i \neq j$), resulting in an SINR of ρ_j from desired stream j . This optimization problem in fact has a closed-form solution: if the channel matrix is \mathbf{H} , the best precoder is the pseudoinverse of the channel matrix, \mathbf{H}^\dagger . Subsequently splitting power among all data streams keeps the power allocation process completely decoupled from precoding, making ZFBF a highly lightweight, yet efficient solution that is attractive for implementation.

The per-antenna power constraint. Since each antenna has its own power amplifier, 802.11ac specifies a per-antenna power constraint P to replace the total power constraint (2a):

$$\sum_{j=1}^{|\mathcal{C}|} |v_{kj}|^2 \leq P, \forall k \in \mathcal{T} \quad (3)$$

In the presence of this constraint, ZFBF may still distribute power equally across data streams, the sum power from all streams on one antenna might either exceed or fall short of the per-antenna power constraint, since each antenna serves a combination of all the data streams according to the precoding vectors $\{\vec{v}_j\}$. Furthermore, there is no closed-form solution to this challenging optimization problem [28], and employing numerical techniques incurs significant complexity without guaranteeing to converge to an optimal solution [11, 32]. Hence, the key challenge is to design a precoding solution that is as lightweight as ZFBF for implementation purposes but at the same time performs an intelligent power allocation across the data streams to leverage the available per-antenna power effectively and thus maximize performance.

Relating Precoding to SINR. Let $\mathbf{V} = [\vec{v}_1 \cdots \vec{v}_{|\mathcal{C}|}]$ be the precoding matrix applied at the AP and \mathbf{S} be the matrix of resulting SINR values at the clients. Entry $s_{ij} \in \mathbf{S}$ measures the power from stream i received at client j , and so captures the interference power from stream i at client j when $i \neq j$. The SINR ρ_j for the desired stream at client j is

$$\rho_j = \frac{s_{jj}}{1 + \sum_{i:i \neq j} s_{ij}} \quad \text{where } s_{ij} = \frac{|\sum_{k=1}^{|\mathcal{T}|} h_{jk} v_{ki}|^2}{N_o} \quad (4)$$

¹The notation $(\cdot)^T$ indicates transpose.




Figure 3: Capacity drop due to naïve power scaling.

Since interference is completely nulled in the case of ZFBF, the SINR matrix reduces to a diagonal SNR matrix.

While ZFBF equally distributes total power among the streams it may not satisfy the per-antenna power constraint (3). One might wonder if a simple power reduction of individual streams on just the antennas that violate (3) would solve the problem. But note that under such a scaling, S becomes non-diagonal and inter-stream interference arises, significantly reducing performance.

Naïve power scaling. To preserve the ZFBF property, we might apply a “global” power reduction, scaling all streams on all antennas by the same value. This means that we would find the antenna k^* that violates (3) by the greatest amount:

$$k^* = \arg \max_{k: \sum_j |v_{kj}|^2 > P} \sum_j |v_{kj}|^2 \quad (5)$$

and then scale all streams on all antennas by a factor of $1/\sum_j |v_{kj}|^2$ so that (3) holds. Such an approach has the potential drawback that the available power on some of the antennas may be significantly underutilized in the process. While this may not be a big problem for CAS, the inherent imbalance in the topology of the antennas and clients in DAS results in a large deviation between the maximum and minimum power allocation to antennas. To illustrate this effect, Figure 3 shows the capacity drop with this naive extension of conventional ZFBF solution to meet the per-antenna power constraint. This trace-based simulation result (for details of testbed, see §5) is for a setup with single four-antenna AP and four single-antenna clients, and clearly indicates that the sub-optimality of a naive power allocation scheme is much more in DAS than in CAS.

3.1.2 Power-balanced precoding

The challenge in designing a lightweight yet efficient precoder for MIDAS is to address the tension between the need to provide interference-free streams and the need to deliver higher power to clients under the per-antenna power constraint. MIDAS balances both these needs to deliver improved MU-MIMO performance. Note that after ZFBF is applied, it is necessary to scale down certain rows of the precoding matrix (corresponding to antennas) that violate the per-antenna power constraint, but this raises the two following problems. Firstly, while different elements in the row may be scaled by different amounts, this must be carefully done as scaling each element of the row directly reduces the SNR and hence rate of one of the streams. Secondly, it is important to scale the elements in the same column (corresponding to a client) by the same factor in order to retain the interference-free property of ZFBF. Hence, the challenge is to pick *scaling weights* w_j ($0 < w_j \leq 1$) for each client j to apply to the ZFBF precoding solution that will ensure all rows satisfy the per-antenna power constraint, while at the same time keeping the associated loss in power and hence in rate to a

$$\underbrace{\begin{bmatrix} w_1 v_{11} & v_{12} & v_{13} & v_{14} \\ w_1 v_{21} & v_{22} & v_{23} & v_{24} \\ w_1 v_{31} & v_{32} & v_{33} & v_{34} \\ w_1 v_{41} & v_{42} & v_{43} & v_{44} \end{bmatrix}}_{\text{Scaled precoding matrix } \mathbf{V}} \rightarrow \underbrace{\begin{bmatrix} w_1^2 s_{11} & 0 & 0 & 0 \\ 0 & s_{22} & 0 & 0 \\ 0 & 0 & s_{33} & 0 \\ 0 & 0 & 0 & s_{44} \end{bmatrix}}_{\text{SINR matrix } \mathbf{S}}$$

Figure 4: SNR impact of scaling a client’s precoding in \mathbf{V} .

minimum. MIDAS realizes this through following novel, iterative, *power-balancing mechanism*:

- **Step 1:** Apply ZFBF to compute the precoding matrix \mathbf{V} .
- **Step 2:** Normalize each of the columns (clients) of \mathbf{V} , applying equal power to each stream.
- **Step 3:** Pick the row (AP antenna) k^* that violates the per-antenna power constraint by the most. Determine the appropriate scaling weight w_j ($0 < w_j \leq 1$) for each element (client) in this row through reverse water-filling to correct the violation.
- **Step 4:** Apply the scaling weight (w_j) computed for each element (v_{k^*j}) of this row (k^*) to the *entire* column corresponding to that stream (*i.e.* $v_{kj} \leftarrow w_j v_{kj}$, $\forall k, j$), thereby retaining the interference-free property of ZFBF. With the updated precoding values, repeat steps 3 and 4 till all rows satisfy the per-antenna power constraint.

Restoring the per-antenna power constraint (Step 3). We now elaborate on the key and novel step of determining the right set of weights w_j to satisfy (3). The received SINR ρ_j in ZFBF is related to the precoded values through (4), where $\rho_j = s_{jj}$. Approximating $\sum_j \log_2(1 + \rho_j) \approx \sum_j \log_2(\rho_j)$ for moderate to large ρ_j , we find that when we scale a column of \mathbf{V} by w_j , the net rate decreases to

$$\sum_j \log_2 \left(\frac{|\sum_{k=1}^{|T|} h_{jk} v_{kj}|^2}{N_o} \right) + \sum_j \log_2(w_j^2) \quad (6)$$

A 4×4 MU-MIMO precoding matrix scaling is shown in Figure 4. Note that while the rate reduction due to a scaling factor w_j is constant ($\log_2(w_j^2)$) and independent of the precoding value to which it was applied, the corresponding reduction in power allocated to different streams in the same antenna ($|w_j v_{kj}|^2$) does depend on that precoding value. *Scaling larger precoding values thus contributes to a larger reduction in power, thereby helping meet the power constraint on the antenna with less loss in rate.* MIDAS leverages this key observation in its precoding power allocation algorithm.

When the same scaling factor is applied to all the elements of a row (antenna), this does not leverage the large variation in the precoding values of the row. Given that the latter is typical of a DAS deployment, this results in a large loss in rate. Hence, MIDAS determines a *separate* scaling factor for each of the elements in the row, taking into account their precoding values. It adopts an approach similar in principle to the classic technique of *waterfilling* [27], albeit tailored for our problem. In conventional waterfilling, given some power values in different bins and a remaining power budget, the goal is to find power allocations to the bins such that the aggregate rate is maximized. We have two requirements specifically for our problem: (i) we need to potentially re-allocate power in *all* the rows and zero power allocation to any of the streams is not allowed as it will eliminate the use of that stream all together, and (ii) when replacing existing power allocations in a row, we don’t allow a stream to receive more power than its current allocation because this allocation increase will apply to the same stream in other rows as well, thereby resulting in potential power violations in other rows that were previously satisfied. This is particularly serious




Figure 5: Reverse water-filling concept.

as it can result in oscillations and getting stuck in an infinite loop, preventing convergence to a solution.

Reverse Waterfilling. For these reasons, we address the reverse version of the waterfilling problem, wherein we compute the power reductions for each stream in the row so as to satisfy the corresponding antenna power constraint, while keeping the rate loss to a minimum. We refer to this problem as *reverse waterfilling*. Let k be the row in which power violation needs to be restored. The goal is to now maximize the net rate, while ensuring that the power reductions P_j (resulting from down-scaling the precoding values, i.e. $w_j v_{kj}$) restore the per-antenna power constraint.

$$\begin{aligned} & \text{Maximize } \sum_j \log_2 \left(1 + w_j^2 \rho_j \right) \\ & \text{subject to, } \sum_j (|v_{kj}|^2 - P_j) \leq P \end{aligned} \quad (7)$$

where $w_j^2 = 1 - \frac{P_j}{|v_{kj}|^2}$. The Lagrangian for the above problem with respect to the vector of power reductions \mathbb{P} is

$$L(\mathbb{P}) = \sum_j \log_2 \left(1 + \left(1 - \frac{P_j}{|v_{kj}|^2} \right) \rho_j \right) - \lambda \left(\sum_j (|v_{kj}|^2 - P_j) \right) \quad (8)$$

Applying KKT conditions, one can obtain the power reductions as

$$P_j = \left[\left(1 + \frac{1}{\rho_j} \right) |v_{kj}|^2 - \frac{1}{\lambda} \right]^+ \quad (9)$$

where $[x]^+ = \max\{x, 0\}$. We now compute λ by applying the $\{P_j\}$ in the power constraint in Eqn. 7. The resulting scaling weights are $w_j = \sqrt{\frac{1}{(|v_{kj}|^2 * \lambda) - \frac{1}{\rho_j}}}$. An illustration of our reverse waterfilling solution is shown in Figure 5. Note that by computing only positive power reductions, we ensure that prior rows for which power violations have been restored are not affected (thereby ensuring convergence). Zero power allocation is not allowed so none of the streams are removed from transmission.

Knowledge of the precoding values allows our solution to leverage the topology imbalance in DAS, while keeping the rate loss to a minimum. This results in a large improvement over applying conventional ZFBF-based MU-MIMO solutions for DAS (see Section 5). Further, computing the reverse waterfilling solution is a lightweight operation, can be done in closed form and needs to be run at most $|T|$ times. This makes the proposed power-balanced precoding highly amenable to real-time implementation.

3.2 DAS-aware MAC for MU-MIMO

We first provide an overview of the joint MAC and PHY operations in MIDAS. Then, we discuss its various design elements that handle the MAC challenges unique to executing MU-MIMO over a DAS deployment in 802.11ac.

3.2.1 Overview

In MIDAS, each of the antennas at an AP competes for access to the channel independently. The sequence of MU-MIMO operations at each MIDAS AP is as follows:

Step 1: Opportunistic Antenna Selection - Once an antenna at an AP gains channel access, the AP checks the channel status (network allocation vector, NAV timers) of other antennas to see if more antennas can be opportunistically added (§3.2.2, §3.2.3).

Step 2: Virtual Packet Tagging - Each of the packets in the AP's queues are tagged with a subset of the DAS antenna IDs as their preferred antennas (based on channel strength) for transmission to the respective clients (§3.2.4).

Step 3: Client Selection - For each of the available antennas (say n antennas), the AP determines the packet (client) that needs to be scheduled based on both its tag as well as its fairness counter. The antennas are considered in order of their NAV timer expiration. Each antenna chooses a different client, resulting in n potential clients (§3.2.5) for MU-MIMO transmission.

Step 4: Channel Estimation - Once a subset of n clients are determined for MU-MIMO, the AP initiates the channel estimation procedure and obtains the channel state information necessary for precoding (§3.3).

Step 5: Power-balanced Precoding - The AP applies the power-balanced precoding solution to execute MU-MIMO from the available antennas to the selected clients (§3.1).

Step 6: Counter Updates: The fairness counters for all the clients serviced are updated and the procedure repeats.

3.2.2 Channel access at antenna granularity

Challenge: While the ability to sense the channels around the antennas independently allows for finer spatial reuse with DAS, the legacy MAC design in 802.11ac cannot leverage MIDAS's full potential for spatial reuse. This is because 802.11ac's MAC is still designed from the perspective of a CAS, allowing for a single channel state at the transmitter. The approximation of the channel states at multiple antennas with a single channel state parameter becomes highly sub-optimal in a DAS set-up, where the antennas could sense highly disparate channel states (some being idle while others are busy). Engaging all the DAS antennas using the single channel state conservatively (busy even if one antenna busy) would prevent us from leveraging the fine grained spatial reuse offered by DAS, while using it aggressively (idle even if one antenna idle) would exacerbate the hidden terminal problem in MU-MIMO. The key challenge here is to instrument 802.11ac's MAC to allow for independent channel sensing at each of the antennas and leverage it effectively to improve spatial reuse with DAS.

Virtual Carrier Sensing Per-antenna: In addition to physical carrier sensing at each of the antennas, MIDAS allows for independent virtual carrier sensing at each of the antennas by provisioning as many network allocation vector (NAV) timers as antennas. Each of these NAVs captures the duration for which the channel around an antenna is busy (based on overhead data or RTS/CTS headers), thereby indicating which antennas are actually available for a MU-MIMO transmission at any given instant. As we will discuss below in Section 3.3, none of the other MAC related parameters in 802.11ac such as contention window, backoff mechanisms, etc. need to change, making the modifications easier to effect.

3.2.3 Opportunistic antenna selection

Challenge: Since each distributed antenna can sense different channel states at the same time, it is possible that while one antenna is available for transmission, another antenna may not. However, the busy antenna may soon become available within a short dura-

tion in the future (based on its NAV timer). An important question arises here: *is it better to wait for additional antennas to become available for MU-MIMO transmission or go with current available antennas?* While waiting for more antennas to become available will increase the number of streams that can be multiplexed in MU-MIMO, thereby delivering better throughput, it could also waste precious air-time in waiting and hence lose throughput.

Opportunistic Selection: To strike a balance, MIDAS *opportunistically* picks antennas for MU-MIMO transmission. When one of the antennas gain access to the channel, the AP checks the NAV timers of the remaining antennas to see if any of them will expire within a DIFS duration. If so, it will wait for those particular antennas to become available and use them all towards MU-MIMO transmission. MIDAS selects the DIFS to be its short waiting period to accumulate antennas for MU-MIMO transmission for the following reasons: (i) DIFS is the basic time unit that every AP needs to check for an idle channel before proceeding with any transmission, (ii) it gives reasonable room to accumulate more antennas for MU-MIMO transmission, and (iii) is not too long to lose out channel access that has been gained.

3.2.4 Filtering clients based on available antennas

Challenge: Unlike CAS, the large disparity in the channel gains from different antennas to a client in a DAS, results in some antennas (with high gains) being naturally preferred for transmission to a client over other antennas. Further, when an antenna's channel state is busy, serving a client in its vicinity from another antenna that is free but farther away, is highly in-efficient. First, the channel gain and hence rate to the client from the farther antenna would be low. More importantly, the channel state of the antenna close to the client being busy should reflect the potential state of the client as well. Hence, making transmissions to such a client might not just be futile but might add to hidden terminals in the network.

Virtual Packet Tagging: Based on the average received signal strength from the different antennas at each client, the MIDAS AP orders the antennas in a decreasing order of preference with respect to each client. It then (virtually) tags each packet meant for a client with the two best antennas from the client's preference list. In other words, a packet for this client is considered for MU-MIMO transmission only if at least one of the two antennas it has been tagged to is available as illustrated in Figure 6. We can see that while A3 and A4 are busy, client 5 and 6 will not be tagged for transmission this round. This not only provides higher transmission rates but also leverages the independent channel sensing feature at antennas effectively to alleviate hidden terminals. Tagging only a single antenna per client can result in under-utilization in some scenarios (*e.g.* clustered user distribution), where some antennas, albeit available, are not leveraged since they have not been tagged by any packets in the AP's queue. Similarly, tagging all antennas per client can result in inefficient topologies as antennas will choose the client far away. With a medium client density, MIDAS tags *two* antennas per client to ensure efficient utilization of AP's antenna resources. This number can be adjusted accordingly with different client densities. Note that the goal of packet tagging is to aid client selection by filtering out clients that are inappropriate for communication from a given set of available antennas. Appropriate selection of clients leads to higher overall capacity.

3.2.5 Antenna-specific client selection

Challenges: An important challenge closely tied to the MAC design is the selection of clients for MU-MIMO transmissions. Once each antenna has been tagged with clients, we need to select a set of clients to receive packets. It is well known that the net rate of a




Figure 6: Virtual packet tagging enables more efficient precoding which leads to higher capacity.

MU-MIMO transmission depends on this set of clients chosen for the transmission. While client selection is important even for MU-MIMO in CAS, its significance is amplified in DAS. Due to the distributed topology of antennas, a large difference in the channel gains from different antennas to a client in DAS (unlike in CAS) can be leveraged to improve MU-MIMO performance. It can also hurt performance if clients are not properly chosen. Hence, picking the right set of clients for MU-MIMO has a much larger impact on MU-MIMO performance in DAS than in CAS. Further, client selection policies are implementation dependent and are typically a function of either just fairness or both fairness and instantaneous rate. Incorporating the latter for MU-MIMO is challenging because selecting the best set of clients for a given MU-MIMO transmission requires estimation and feedback of *all* clients' channels to make the decision. However, *measuring CSI from all clients involves not just the overhead of a CSI request (from AP) and response (from all clients), but also significant latency before being able to determine the best subset of clients.* The latter potentially making the estimates too stale to be useful in the first place due to small channel coherence time indoors. One might wonder if prior channel estimates can be used to perform client selection. For this to work, the prior estimates must still be within the coherence time of the clients' channels, which is in the order of tens of milliseconds for a day-time environment in enterprises, etc. [27].

Antenna-specific, Fairness-driven Selection: MIDAS leverages the DAS topology to select clients purely based on their antenna preference and fairness, thereby eliminating the associated channel estimation overhead. Based on opportunistic antenna selection, MIDAS orders the antennas according to their NAV expiry (earliest being first). It then picks the first antenna (*primary antenna*) that gained channel access and applies the scheduling policy only among those packets in the queue that have been tagged (preferred) to this antenna and picks a single packet. It then moves to the next antenna (*secondary antennas*) in the list of available antennas for MU-MIMO transmission and repeats this process to determine packets for all available antennas. Here, since a client's packet can tag two antennas, a client selected for a prior antenna is excluded from being re-considered for a later antenna. Once the set of clients is selected, MU-MIMO transmission is executed jointly from the available set of antennas using the precoding algorithm proposed. *Note that although one client is selected for each antenna, the data streams are transmitted from all the antennas to all the clients with precoding rather than each antenna transmitting to one client.*

Also note that in CAS, since all antennas are similar to each client, CSI estimation and scheduling are tightly coupled and all the clients needs to be considered for the best subset for MU-MIMO. However, due to the topology imbalance in DAS, *picking clients based on their antenna preference contributes to a large part of the client selection gain.* This in turn allows MIDAS to decouple CSI estimation from scheduling without an appreciable loss in MU-MIMO performance.

Scheduling Policy: In terms of the scheduling policy itself, one can employ any desired scheduling policy ranging from simple round-robin to more sophisticated weighted fair queuing schemes with delay guarantees in MIDAS. In particular, deficit round robin (DRR) is a type of weighted fair queuing that also incurs low complexity and is hence conducive for implementation. MIDAS employs DRR as its scheduling policy, albeit tailored for MU-MIMO, but can easily accommodate other scheduling policies as well. Briefly, a deficit counter is maintained for each client (in terms of number of time slots), indicating a measure of pending service for the client for fairness purposes. With each transmission opportunity (called TXOP in 802.11) spanning a contiguous set of time slots (few milliseconds, say T), once the primary antenna gains channel access, the client with the largest deficit counter among those tagged for that antenna is chosen for transmission. Similarly, clients for the other secondary antennas are also chosen based on their deficit counters. Then the deficit counters for all the chosen clients in MU-MIMO transmission are decremented by T , while those that had outstanding packets but did not get selected are incremented by $\frac{nT}{m}$ each, where n is the number of antennas participating in MU-MIMO transmission and m is the number of active clients not chosen for the transmission. In other words, the current service (nT) is distributed equally among the non-scheduled clients to guide the schedule towards a fair allocation.

3.3 Realizing MIDAS within 802.11 ac

MU-MIMO Precoding: 802.11ac provides the framework (channel estimation, precoding, etc.) to realize MU-MIMO. The determination of the specific precoder to apply is implementation dependent. However, since each antenna has its own power amplifier, the precoder needs to satisfy the per-antenna power constraint. Given that the power-balanced precoding in MIDAS is specifically designed to address the per-antenna power constraint and leverage the DAS deployment, it can be easily realized with 802.11ac.

MU-MIMO MAC: 802.11ac adopts a version of 802.11e’s MAC and re-purposes it for MU-MIMO. 802.11e was introduced to provide quality of service (QoS) and service differentiation between different traffic classes. It employs four different transmission queues to account for four traffic classes, namely video, voice, best effort and background traffic. Different contention window and timing values are used for the different classes to explicitly prioritize between the classes in their access mechanism. Each of these traffic classes contend internally at the AP, while the AP contends externally with other APs. When the AP gains channel access, only one of the traffic classes gains access. 802.11ac re-purposes these four different queues to enable MU-MIMO transmissions. When one of the traffic classes gains channel access, it is made the primary access class. If sufficient clients (data) are not available in the primary class, clients from secondary traffic classes can be chosen.

MIDAS’s MAC operations were explained with respect to a single traffic class. Being designed with 802.11ac in mind, it is easy to see it will readily fit in 802.11ac’s MAC. The explicit priorities and channel access procedure for the traffic classes will determine the primary access class as before. Once the primary and secondary traffic classes are determined, MIDAS’s client selection and scheduling procedure can be readily applied within each of the traffic classes.

Channel Estimation: Finally, 802.11ac provides complete and elaborate support for channel estimation (through a process called *sounding* with the help of null data packets) and feedback that are critical for MU-MIMO precoding. However, such estimation is invoked only after an AP gains channel access and hence *only CSI of those clients chosen for MU-MIMO transmission will be available to the AP*. Note that, MIDAS employs an antenna-specific, fairness-

based client selection process that selects efficient client topologies for MU-MIMO without relying on CSI estimates. This makes it well-aligned with the support provided by 802.11ac.

4. IMPLEMENTATION

Testbed and Deployment: Our goal is to realize MIDAS with off-the-shelf 802.11ac APs and clients. However, the current 802.11ac APs do not implement MU-MIMO yet, which is expected to be a part of the next release, scheduled for the end of this year. Hence, we implement MIDAS on the WARP [20] platform using a representative version of the relevant PHY and MAC mechanisms in 802.11ac that we need for our purpose. Each WARP board can support up to four RF front-ends and is connected to a host PC. The host PC performs the (de-) construction and processing of baseband waveforms to and from the FPGA board with the RF front-ends. The WARP board along with the host PC serves as the AP itself, with its four RF boards connected with four antennas. Single antenna WARP boards are used for the clients. To realize a DAS set-up, the four antennas are distributed around the AP with the help of RF coaxial cables.

PHY/MAC Implementation: All the baseband processing happens on the main board. With the host PC processing the transmitted (received) data, our power-balanced precoding mechanism is implemented in the host PC and applied to the data streams that are transmitted through the RF front-ends. The precoding is performed while adhering to the constraints imposed by 802.11ac’s PHY specification. With each of the RF front-ends having their own power amplifiers, this incorporates the per-antenna power constraint imposed by the 802.11ac specification. We significantly instrument the basic CSMA MAC in WARP to realize the DAS-aware MAC in MIDAS. We start with a basic implementation of 802.11g for WARP and then incorporate the relevant MAC features from 802.11ac that are needed for our purpose. Specifically, we provision individual MAC states, *i.e.* the virtual carrier sensing parameter (network allocation vector, NAV) and enable packet detection for each of the DAS antennas. We implement multiple NAV states, individual packet detection and opportunistic antenna selection along with the notion of virtual packet tagging. We also implement the deficit round robin policy for scheduling multiple clients tagged to an antenna. Because of the large delay of WARP platform, we are still not able to realize a close-loop MAC. Thus we implement the MAC schemes on top of the latest WARP 802.11 Reference Design with Xilinx XPS 13.4 and feed the MAC layer decision into the PHY layer implementation, in which we implement the channel estimation, power-balanced precoding and upto 4x4 MU-MIMO on WARPLab reference design.

5. EVALUATION

We break down our experimental evaluation of MIDAS into three parts: (i) PHY layer precoding, (ii) MAC layer features, and (iii) end-end PHY+MAC. We also evaluate MIDAS in larger topologies through simulations using channel traces.

5.1 Methodology, Baseline & Metrics

We deploy our testbed in both indoor enterprise and academic lab environments. We consider up to three APs, each with at most four antennas and capable of performing upto 4×4 MU-MIMO. APs are deployed with an inter-AP distance of around 15 m, while the WARP client boards are located in offices/rooms as well as in the corridor. The antennas are co-located at the AP in CAS with a half-wavelength distance in-between, while they are distributed around the AP in DAS with a distance of 5 - 10 m from the AP. Each antenna has its own power amplifier in both CAS and DAS. We run our

experiments by sending back to back packets for 10 seconds from the AP to clients and recording the results by averaging over multiple runs for each topology and over 50 to 100 different topologies.

We evaluate the benefits from MIDAS’s precoding and MAC mechanism against several baselines. For precoding, we evaluate our power-balanced precoding against both the optimal precoding through the MATLAB numerical toolbox as well as a simple extension to conventional ZFBF precoding. The latter first splits power equally to all the streams, then scales the power of all the antennas by a common factor in order to satisfy the per-antenna power constraints. For MAC, we evaluate our proposed MAC against a CAS system with a conventional CSMA MAC.

We translate measured SINR for the data streams at different clients into network capacity using the Shannon capacity formula. Note that, unlike for open-loop MIMO schemes, where rate adaptation is needed, for closed-loop MIMO schemes like MU-MIMO, the CSI information is directly used to determine the appropriate rate (modulation and coding) that must be applied for a stream, eliminating the need for an explicit rate adaptation process. Hence, measured SINR from experiments would directly translate to capacity as well.

5.2 Power-balanced Precoding in MIDAS

5.2.1 Benchmarking link gain from DAS

The CDF curves of link layer SNR for CAS and DAS during SISO transmissions from a single AP are shown in Figure 7. The CAS antenna positions are fixed while DAS antennas and clients are randomly deployed in 60 different topologies. Each AP has four antennas attached. The SNR value is measured when only one antenna transmits to the client. Each client is mapped to one antenna in a greedy way: the client with the strongest signal is chosen and this mapped antenna-client pair is excluded for next round mapping. It can be seen that *the shorter DAS links to clients provide a median link gain of 5 dB with four antennas*. This gain keeps increasing with more clients and antennas.

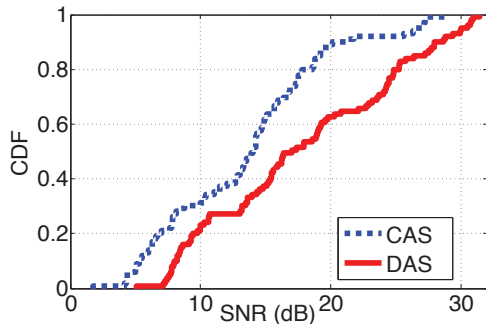


Figure 7: Distribution (across clients) of link layer SNR.

5.2.2 MU-MIMO capacity increase

In order to evaluate the MU-MIMO capacity increase, we fix the CAS antenna positions of a single AP and randomly deploy the clients and DAS antennas. We present the overall capacity (calculated from measured SINR) of MIDAS and CAS with 2x2 MU-MIMO and 4x4 MU-MIMO in Figure 8 and Figure 9. Here, while CAS employs the baseline precoding mechanism, MIDAS employs our power-balanced precoding. We carry out experiments in two different indoor office environments: an enterprise (Office A) and a graduate student lab (Office B), which is more crowded. We can see that *MIDAS has a median gain of 40-67% over CAS for two antennas, that increases to 45-80% with higher number (four) of*

antennas. The gains are contributed both by the link gain from DAS as well as the power-balanced precoding in MIDAS.

5.2.3 Impact of power-balanced precoding

We isolate the impact of our smart, power-balanced precoding scheme in Figure 10 for the four antenna case (Office B). We compare the performance of DAS and CAS when applying our power-balanced precoding respectively. While our *smart precoding also helps CAS improve its median capacity by 12% (11.5 to 13 bits/s/Hz), the improvement is much greater for DAS (about 30%)*. Thus, the benefits of the power-balanced precoding are more pronounced for DAS than for CAS. This also corroborates our intuition that a naive power allocation scheme (like the baseline), albeit fine for CAS, is not sufficient for DAS due to its inherent topology imbalance.

We also compare our smart precoding scheme and the optimal precoder obtained from solving the precoding optimization problem through the MATLAB numerical solver/toolbox in Figure 11. The evaluation is conducted on both the testbed as well as through traces. In trace-based simulations, we see that our *precoding in MIDAS performs efficiently within 99% of the optimal, while not incurring the complexity or processing latency of the latter*. In the experiments, interestingly, in a few of the topologies, the MIDAS precoding scheme even outperforms the optimal precoder. One possible reason for this is that it takes a couple of seconds for the optimization toolbox to calculate the optimal precoder. By the time the precoder is determined, the channel could have changed (even slightly), resulting in the precoder being no longer optimal. This also reinforces our claim that *a fast, light weight precoding algorithm is essential for MU-MIMO given the short coherence time of the wireless channel*.




Figure 8: MIDAS capacity increase (Office A).

5.3 DAS-aware MAC in MIDAS

We evaluate in this section MIDAS’ MAC described in §3.2.




Figure 9: MIDAS capacity increase (Office B).

5.3.1 Number of Simultaneous Transmissions

We first evaluate the total number of simultaneous transmissions (spatial reuse) that can be supported by the MAC. We deploy three APs that can overhear each other. For CAS setup, one AP can be activated at a time to support four simultaneous transmissions. As the antennas are distributed at different locations in MIDAS, the NAV values and channel sensing results are different at each DAS antenna. The channel is thus utilized more efficiently in a finer manner. To evaluate the number of simultaneous transmissions can be supported at the MIDAS APs, we randomly enable one to four transmissions at the first AP *A* and then check the number of transmissions can be supported simultaneously at the second AP *B* based on its NAVs and channel sensing results. Next, we enable all the transmissions supported at both *A* and *B* and similarly evaluate the number of transmissions that can be simultaneously supported at AP *C*. Note that more simultaneous transmissions does not necessarily mean higher overall capacity. The overall capacity also depends on the capacity of each transmission. We show the end-to-end capacity evaluation in Section 5.4. We run experiments for 30 topologies generated by random antenna deployments. In order to avoid bias in results, for each deployment, the following requirement is satisfied: any two antennas from the same AP cannot be deployed within a 60-degree sector measured with respect to the AP. This ensures that antennas are not clustered at the far end with respect to the other two APs to favor the results. Our testbed results in Figure 12 show that among all the 30 topologies, only two MIDAS deployments support less transmissions compared to CAS. *The median improvement in the total number of transmissions is around 50%, while 90% improvement can be achieved for some topologies. This can be directly attributed to the fine grained spatial reuse enabled in MIDAS to leverage the DAS deployment. Note that with the same number of transmissions in Section 5.2.2, MIDAS already outperforms CAS.*




Figure 10: Impact of Smart Precoding in CAS and DAS.

5.3.2 Significance of virtual packet tagging

We consider a scenario where two antennas (out of four) at a MIDAS AP are available at the MAC layer, while there are a total of four clients to be selected for transmission. With virtual packet tagging, MIDAS chooses a more appropriate group of two clients. We also assume there are always packets for all the four clients in the queue. We present the results in Figure 14 and compare this scheme against a scheme that chooses two clients randomly. We can clearly observe a *50% increase in median capacity that results from picking clients intelligently to aid in better MU-MIMO transmissions.*

5.3.3 Deadzone reduction

We now show that a MIDAS deployment can help alleviate the problem of signal deadzones. We deploy the AP in both MIDAS and CAS modes, with the DAS antennas randomly deployed surrounding the AP. Then we carry out measurements within the coverage area for deadspots and one measurement is taken every 0.5 m. We show one typical MIDAS/CAS deadzone comparison map in Figure 13. The gray small squares indicate the deadspots identified. We can clearly see that deadzone spots are much fewer in MIDAS. *On average of 10 different random antenna deployments, DAS deployment helps provide more effective client coverage, resulting in a significant reduction of deadspots by about 91% compared to CAS.*

5.3.4 Alleviating hidden terminals

We setup two APs, AP1 and AP2. They are separated far enough that they cannot overhear each other but not too far to eliminate hidden terminals. We distribute the DAS antennas randomly around APs within a distance of 50%-75% of the CAS AP's transmission range from the original AP. We then measure the number of hidden terminal spots in the area with a distance of one meter between each measured spots in 10 random antenna deployment topologies. We found that, *on average DAS has 94% hidden terminal spots*




Figure 11: Comparing MIDAS vs optimal precoding.




Figure 12: Ratio of simultaneous transmissions.

removed compared with CAS, thereby alleviating hidden terminals significantly. This can be attributed to two reasons: (i) distributed antennas in MIDAS are able to sense a larger region compared to collocated antennas in CAS, and (ii) while the transmission power is maximal at the AP and decreases with distance in CAS, it is more evenly distributed around the AP in DAS, allowing for better detection by other DAS/CAS APs.

5.4 End-End Evaluation of MIDAS

The setup consists of a three-AP network. Each of the APs is capable of up to 4×4 MU-MIMO transmissions. Each AP has four randomly deployed clients around it. The CAS APs employ the traditional CSMA MAC with the baseline precoding scheme while the MIDAS APs employ the DAS-aware MAC coupled with our power-balanced precoding described in Sections 3.2 and 3.1 respectively. The overall capacity performance between CAS and MIDAS is shown in Figure 15. The CDF is measured over 60 different topologies. MIDAS achieves a significant improvement of



Figure 13: Comparing deadzones in MIDAS and CAS

200% in capacity gain over CAS. The substantial PHY+MAC gain in MIDAS is the culmination of three factors: (i) For the same number of streams, MIDAS outperforms CAS in MU-MIMO performance; (ii) For the same number of APs, more simultaneous streams are enabled in MIDAS due to its DAS-aware channel-sensing MAC; and (iii) Virtual packet tagging in MIDAS takes advantage of the DAS topology to further increase MU-MIMO performance through intelligent client selection.

5.5 Large Scale Trace-based Simulations

We carried out trace-based simulation to verify MIDAS end-to-end performance in a large scale setup. We randomly deploy eight APs in a 60×60 m region. All the APs are capable of performing up to 4×4 MU-MIMO transmission. We make sure none of the CAS APs can overhear more than three other APs for each deployment. For the DAS antenna deployment, none of the antennas can be deployed out of the original AP coverage area and no two antennas can be deployed within 5 m of distance. We measure the CSI and the interference level between AP antennas and clients and feedback the traces to the simulation. The overall result is shown in Figure 16. We can see that the large scale simulation result follows the trend we obtain in testbed: DAS outperforms CAS by more than 150%.

6. RELATED WORK

MIDAS makes contributions in the areas of MU-MIMO precoding power allocation and medium access control. In this section we survey the related work in each area and place MIDAS's contributions into context. We close with a review of other DAS systems and theoretical proposals.

MU-MIMO precoding. Several systems have taken multi-user MIMO [3, 26], interference alignment [8, 15] and distributed MU-MIMO [19, 5, 13, 30] from theory to practice. Falling under the multi-AP deployment model discussed in Section 2, their precoding solutions are neither designed to leverage the topology imbalance in DAS, nor satisfy the per-antenna power constraint imposed by 802.11ac in a lightweight manner. Other precoders specifically designed for DAS systems [14] are designed only for a sum-power constraint (§3.1) while [25] is designed for large number of users and high mobility users. On the other hand, prior precoders designed with the per-antenna power constraint in mind are unfortunately too complex to realize in practice [11, 18, 32] (and references therein). Waterfilling is a popular technique used to allocate a fixed power budget to orthogonal channels of varying gains, such as OFDM subcarriers. It has been extensively studied, both for Gaussian [27] as well as arbitrary input distributions for SISO and MIMO systems [17]. Prior work in power allocation for DAS systems either focus on




Figure 14: Virtual packet tagging effect




Figure 15: 3-AP testbed




Figure 16: 8-AP simulation.

single-user communications [12, 31], or operates under a sum power constraint instead of a per-antenna power constraint [9]. The nature of the practical constraints to our MU-MIMO precoding problem over DAS prevents MIDAS from employing these solutions. Being specifically designed for the dual challenges of the per-antenna power constraint and a DAS deployment model, the lightweight precoder advances the state-of-the-art in this area.

Medium access control & rate adaptation. The MAC in MIDAS is also concerned with a virtual selection and tagging of antennas to aid client selection for MU-MIMO. This is different from conventional antenna selection [16], which arises because antennas used for communication are limited by the available RF chains. Recently, NEMOx [32] designed a random access MAC for distributed MU-MIMO systems. However, being designed from scratch, it does not directly lend itself to 802.11ac systems that in turn form our focus. Other work [24] focuses on rate adaption for MU-MIMO.

Other DAS systems. Serving as a simple broadcast medium, DAS has traditionally been used to provide coverage and handle mobility in WiFi and cellular systems [1, 2]. Recent proposals [33, 34] have combined DAS with frequency or spatial reuse to enable multi-user transmissions, albeit still operating under SISO. Other theoretical work [6, 23, 22, 10] has studied the benefits of applying MU-MIMO techniques over DAS, albeit using conventional MIMO precoding solutions that limit their potential to leverage DAS for reasons explained in §3.1.

7. DISCUSSION

Deployment. In most enterprises, APs are attached to walls/ceilings and wired over the ceilings for safety and aesthetics. A similar approach can be adopted in deploying the MIDAS antennas. The RF extension cables can be run over the ceiling, thereby incurring very little additional effort and expense. Deploying DAS antennas too close to the AP makes the system tend to a CAS, leading to reduced channel utilization. On the other hand, positioning too far from the AP may have an adverse impact on links and the MU-MIMO performance. From experience, a good distance is 50-75% of the AP's CAS coverage range. We leave the problem of optimizing placement of antennas open for future work.

Beamforming. While employing multiple antennas to send multiple streams to different clients would leverage multiplexing (MU-MIMO), using all them to alternatively send a single stream to a single client would leverage beamforming. However, beamforming (SNR) gain is logarithmic unlike the multiplexing gain (that is linear). Due to the decoupled nature of the antennas in MIDAS, when all antennas at an AP are leveraged for beamforming, this would not only come at the expense of multiplexing gain but would also silence other antennas (of other APs) in the neighborhood, thereby limiting the potential for improved spatial reuse from MIDAS. Hence, in

low client densities, when an AP is restricted to beamforming to a single client, it is important to leverage only the antennas in the neighborhood of the client for beamforming in MIDAS, thereby not affecting spatial reuse by antennas of other nearby APs.

Co-located antennas in a DAS topology. In this work we deploy DAS antennas separately, but one might wonder whether it would be beneficial to co-locate more antennas, especially if clients have multiple antennas. We argue, however, that deploying non-co-located DAS antennas is still the right approach. If a client had more antennas (say m), and we also deploy m AP antennas at each location, which in turn would reduce the number of DAS deployment locations possible, degrading the DAS deployment advantage. Having more deployment locations (with a single antenna) and allowing them to cooperate still leverage all the multiplexing gain possible, while catering to more clients associated with the AP.

SVD precoding. SVD is employed as a precoding scheme in 802.11n, when multiple streams are sent to a single client with multiple antennas. However, due to the distributed nature of clients in 802.11ac, SVD is not suited because the clients do not share streams with each other in 802.11ac. Even with MIMO transmissions to a single client with multiple antennas in 802.11ac, SVD is still not a good candidate for MIDAS because of the imbalanced topology in DAS and the (equal) power allocation nature of SVD.

8. CONCLUSION

With next generation wireless local area networks moving to multi-user MIMO communication, we argue that the conventional co-located AP antenna topology is inefficient from both physical- and MAC-layer perspectives. Towards helping wireless service providers leverage the most out of 802.11ac APs, we have presented the design and implementation of MIDAS, a multiple input distributed antenna system, that migrates an 802.11ac AP to a DAS topology, supplementing it with a suite of novel PHY and MAC mechanisms to maximize network performance. We have shown that MIDAS's mechanisms are lightweight additions to an AP's software device driver, not requiring any changes to 802.11ac's hardware. We have described our WARP radio implementation and presented results from our experimental testbed that demonstrate MIDAS's techniques will significantly increase the capacity of tomorrow's 802.11ac networks.

9. ACKNOWLEDGEMENTS

The research leading to these results has received funding from the European Research Council under the EU's 7th Framework Programme (FP/2007-2013), ERC grant agreement number 279976. Jie Xiong is supported by a Google European Doctoral Fellowship in Wireless Networking and part of this work was done at NEC Labs America during his internship.

10. REFERENCES

- [1] The hetnet forum. <http://www.thedasforum.org>.
- [2] T. Alade, H. Zhu, and H. Osman. Spectral efficiency analysis of distributed antenna system for in-building wireless communication. In *Personal Indoor and Mobile Radio Communications*, 2011.
- [3] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly. Design and experimental evaluation of multi-user beamforming in wireless lans. In *MobiCom*, 2010.
- [4] Distributed Antenna System: Keeping customers connected, wherever they are (AT&T). <http://www.att.com/gen/press-room?pid=23351>.
- [5] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire. Airsync: Enabling distributed multiuser mimo with full spatial multiplexing. *Networking, ACM Transactions on*, 21(6), 2013.
- [6] B. Cho, S. Seo, and J. Kim. Practical scheme to enable MIMO communications in distributed antenna systems for efficient indoor coverage. In *International Symposium on Communications and Information Technology*, 2009.
- [7] Positioning Statement on Cisco Wireless LAN over Distributed Antenna Systems (Cisco Corp.).
- [8] S. Gollakota, S. D. Perli, and D. Katabi. Interference alignment and cancellation. In *SIGCOMM*, 2009.
- [9] N. Hassan, C. Yuen, and Z. Zhang. Optimal power control and antenna selection for multi-user distributed antenna system with heterogeneous QoS constraints. In *Globecom Workshop*, 2012.
- [10] R. Heath, T. Wu, Y. H. Kwon, and A. C. Soong. Multiuser mimo in distributed antenna systems. In *Signals, Systems and Computers (ASILOMAR), 2010 Conference Record of the Forty Fourth Asilomar Conference on*, pages 1202–1206. IEEE, 2010.
- [11] S. Kaviani and W. Kryzmien. Optimal multiuser zero forcing with per-antenna power constraints for network mimo coordination. *EURASIP J on Wireless Comms. and Netw.*, 2011.
- [12] H. Kim, S.-R. Lee, C. Song, and I. Lee. Optimal power allocation for energy efficiency maximization in distributed antenna systems. In *ICC*, 2013.
- [13] S. Kumar, D. Cifuentes, S. Gollakota, and D. Katabi. Bringing cross-layer MIMO to today’s wireless LANS. In *SIGCOMM*, 2013.
- [14] Q. Li, X. Wang, and X. Feng. Design of linear pre-coders and equalizers for the uplink in multiuser distributed antenna system. In *BMSB*, 2013.
- [15] K. Lin, S. Gollakota, and D. Katabi. Random access heterogeneous MIMO networks. In *SIGCOMM*, 2011.
- [16] S. Mahboob, R. Ruby, and V. Leung. Transmit antenna selection for downlink transmission in a massively distributed antenna system using convex optimization. In *BWCCA*, 2012.
- [17] F. P-Criz, , M. Rodrigues, and S. Verdu. Generalized mercury/waterfilling for multiple-input multiple-output channels. In *ACCC*, 2007.
- [18] Z. Pi. Optimal mimo transmission with per-antenna power constraints. In *IEEE Globecom*, Dec 2012.
- [19] H. Rahul, S. Kumar, and D. Katabi. MegaMIMO: Scaling Wireless Capacity with User Demands. In *ACM SIGCOMM 2012*, Helsinki, Finland, August 2012.
- [20] Rice WARP Platform. <http://warpproject.org>.
- [21] A. Saleh, A. Rustako, and R. Roman. Distributed antennas for indoor radio communications. *IEEE Trans. on Comms.*, COM-35(12):1245–51, 1987.
- [22] S. Schwarz, R. W. Heath Jr, and M. Rupp. Single-user mimo versus multi-user mimo in distributed antenna systems with limited feedback. *EURASIP Journal on Advances in Signal Processing*, 2013(1), 2013.
- [23] C. Shen, S. Zhou, and Y. Yao. Comparison of channel capacity for MIMO-DAS & MIMO-CAS. In *APCC*, 2003.
- [24] W. Shen, Y. Tung, K. Lee, K. Lin, S. Gollakota, D. Katabi, and M. Chen. Rate adaptation for 802.11 multiuser mimo networks. In *MobiCom '12*.
- [25] C. Shepard, N. Anand, and L. Zhong. Practical performance of mu-mimo precoding in many-antenna base stations. In *CellNet '13*.
- [26] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. M. Voelker. SAM: Enabling practical spatial multiple access in wireless LAN. In *MobiCom*, 2009.
- [27] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
- [28] A. Wiesel, Y. C. Eldar, and S. Shamai. Zero-forcing precoding and generalized inverses. *Signal Processing, IEEE Transactions on*, 56(9):4409–4418, 2008.
- [29] H. Xia, A. Herrera, S. Kim, and F. Rico. A CDMA-distributed antenna system for in-building personal communications services. *IEEE JSAC*, 1996.
- [30] Q. Yang, X. Li, H. Yao, J. Fang, K. Tan, W. Hu, Z. J, and Y. Zhang. Bigstation: Enabling scalable real-time signal processing in large mu-mimo systems. In *SIGCOMM '13*.
- [31] X. Zhang, Y. Sun, X. Chen, S. Zhou, J. Wang, and N. Shroff. Distributed power allocation for coordinated multipoint transmissions in distributed antenna systems. *IEEE Trans. on Wireless Comms.*, 12(5), 2013.
- [32] X. Zhang, K. Sundaresan, M. Khojastepour, S. Rangarajan, and K. Shin. NEMOx: Scalable network MIMO for wireless networks. In *MobiCom*, 2013.
- [33] H. Zhu. On frequency reuse in distributed antenna systems. In *VTC*, 2011.
- [34] H. Zhu and J. Wang. Radio resource allocation in multiuser distributed antenna systems. *JSAC*, 2013.