

One Strategy Does Not Serve All: Tailoring Wireless Transmission Strategies to User Profiles

Shailendra Singh*, Karthikeyan Sundaresan†, Amir Khojastepour†, Sampath Rangarajan†, Srikanth V. Krishnamurthy*

*University of California, Riverside {singhs, krish}@cs.ucr.edu
†NEC Laboratories America Inc. {karthiks, amir, sampath}@nec-labs.com

ABSTRACT

The proliferation of smartphones and tablet devices is changing the landscape of user connectivity and data access from predominantly static users to a mix of static and mobile users. While significant advances have been made in wireless transmission strategies (e.g., network MIMO) to meet the increased demand for capacity, such strategies primarily cater to static users. To cope with growing heterogeneity in data access, it is critical to identify and optimize strategies that can cater to users of various profiles to maximize system performance and more importantly, improve users' quality of experience.

Towards this goal, we first show that users can be profiled into *three* distinct categories based on their data access (mobility) and channel coherence characteristics. Then, with real-world experiments, we show that the strategy that best serves users in these categories varies distinctly from one profile to another and belongs to the class of strategies that emphasize either *multiplexing* (eg., netMIMO), *diversity* (eg., distributed antenna systems) or *reuse* (eg., conventional CSMA). Two key challenges remain in translating these inferences to a practical system, namely: (i) how to profile users, and (ii) how to combine strategies to communicate with users of different profiles simultaneously. In addressing these challenges, we present the initial design of *TRINITY* - a practical system that effectively caters to a heterogeneous set of users spanning multiple profiles simultaneously.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication

General Terms

Algorithms, Experimentation, Measurement, Performance

Keywords

Network MIMO, DAS, Reuse, User Profile

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Hotnets '12, October 29–30, 2012, Seattle, WA, USA.
Copyright 2012 ACM 978-1-4503-1776-4/10/12 ...\$10.00.

1. INTRODUCTION

Two trends are becoming evident in enterprise wireless networks. First, enterprises that used to be dominated by static users due to devices like laptops and notebooks are increasingly being populated with mobile users owing to the proliferation of smart phones and BYOD initiatives. Second, wireless transmission strategies have also improved significantly to cope with the increased demand for capacity. Specifically, the past decade has seen strategies go from conventional CSMA to more sophisticated MIMO-based strategies. While standards have come as far as multi-user MIMO (IEEE 802.11ac), research has advanced to the point of building network (distributed) MIMO systems [6, 3].

Network MIMO (netMIMO) allows for multiple data streams to be transmitted concurrently from distributed transmitters to users by converting interference into a multiplexing gain via transmitter cooperation. However, its benefits are realizable only if the whole process of channel estimation, feedback and use for netMIMO operation, is completed within the coherence time of the users' channels. Indeed, applying netMIMO for mobile users or for whom the latter condition is not satisfied, can even hurt performance by as much as 71.4% for a 3 client system as shown later. Given the mix of users with diverse channel and mobility characteristics (referred to as *profiles*) in next generation enterprise networks, it is both important and timely to understand how transmission strategies must be tailored to user profiles so as to maximize system performance and improve user quality of experience.

A potential approach could be to differentiate between static and mobile users, apply netMIMO only for static users, while employing conventional CSMA approaches as in current enterprise networks for mobile users. However, such an approach has two drawbacks: (i) netMIMO could hurt performance even for static users if their channel fluctuations (e.g., environment changes) result in small coherence times; and (ii) while conventional CSMA would work for mobile users, it could be highly sub-optimal due to handover delays and the lack of transmitter cooperation gain.

We argue that users in an enterprise network can be profiled more generally into one of three distinct categories. First, they can be categorized into those with large and small channel coherence times. Further, among those with small coherence times, we can further classify them based on the contributing factor - user mobility or environment dynamics. Similarly, the gamut of wireless transmission strate-

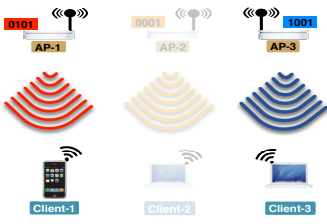


Figure 1: CSMA: (Two slot schedule: AP1,AP3 active simultaneously in slot 1, AP2 in slot 2.)

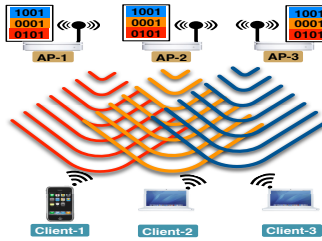


Figure 2: netMIMO: (3 APs sending 3 data streams to 3 clients in each slot.)

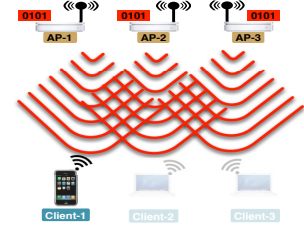


Figure 3: DAS: (3 slot schedule: 3 APs sending same data to one client in each slot.)

gies can be grouped into one of three distinct categories. At the top level, we have those that allow transmitter cooperation and those that do not. We refer to the latter as *reuse* strategies, an example being the current enterprise CSMA scheme, where transmitters reuse the channel based on interference avoidance¹. Schemes that enable transmitter cooperation can be further classified into two categories - *multiplexing*: those that exploit interference (through cooperation) and channel state information (CSI) to transmit multiple independent streams to users (e.g., netMIMO), and *diversity*: those that bypass interference by sending multiple versions of a stream through different transmitters without the need for CSI feedback to provide a diversity gain (e.g., space-time block codes, Distributed Antenna System, etc.)

With the help of implementation of the various strategies and real-world experiments with users of various profiles, we show that the strategy that best serves users of a profile varies distinctly from one profile to another and there is no single strategy that can cater to all users effectively. For example, for a 3 transmitter system, we find that multiplexing schemes (netMIMO) yield the best performance for static users with large coherence times with gains as high as 69.84%. However, when applied to mobile users and even static users with small coherence times, it degrades performance by as much as 71.4%. On the other hand, diversity schemes, with their increased coverage and diversity gain, yield the best performance for mobile users with gains as much as 96.7%. Interestingly, for the third category of static users with small coherence times, transmitter cooperation either degrades performance or results in capacity underutilization, thereby resulting in conventional reuse schemes being the best suited strategy with gains as much as 83.8%.

While our findings would generalize to larger topologies, translating them to a practical system encounters two key challenges: (i) accurate categorization of users into various profiles is central to the whole process, and (ii) how to combine multiple strategies to effectively cater to users of various profiles simultaneously. We present the initial design of *TRINITY* - a practical system to address the aforementioned challenges. Briefly, *TRINITY* is deployed at the central controller (managing enterprise networks) and incorporates three key design elements.

1. It enables simultaneous operation of all strategies by mul-

¹Note that this does not preclude each of the transmitters from individually employing MIMO to their users.

tiplexing them in the frequency domain in OFDM networks, where the available sub-carriers are split between different strategies. This allows for power pooling benefits that are not available with time domain multiplexing.

2. The resources (e.g., # sub-carriers) allocated to a strategy depends on the traffic load of the corresponding user profile and is closely integrated with the user categorization process itself. A measurement based approach coupled with sensor hints (eg., accelerometer, [7]) is employed to accurately categorize users into profiles.

3. To maximize network performance, the performance-complexity tradeoff with multiplexing and the coverage-capacity tradeoff with diversity schemes are further optimized on the sub-carriers allocated to the corresponding user profiles.

2. OVERVIEW OF STRATEGIES

The gamut of transmission strategies can be categorized based on the level of cooperation between the transmitters at the top level. Note that, in all the schemes we discuss below, each transmitter (AP) can individually employ single-user or multi-user MIMO with its own clients. However, for ease of exposition, we will restrict our discussion to single antenna APs and users.

2.1 Non-cooperating Transmitters

The case of non-cooperating transmitters would correspond to the conventional CSMA paradigm - APs avoid interference in the presence of a carrier (sensing), while spatial reuse is automatically leveraged otherwise. In the example in Fig. 1, APs 1 and 3 can transmit in tandem on the same channel, while AP 2 time shares the medium with 1 and 3. Such schemes and their variants (e.g., with MIMO APs) will collectively be referred to as *reuse* schemes.

2.2 Cooperating Transmitters

When transmitters are allowed to cooperate and are synchronized (at symbol level), additional benefits can be leveraged depending on the nature of cooperation.

Diversity: The goal of these schemes is to send multiple, dependent versions of a data stream through multiple transmitters to provide transmit diversity. A popular example would be the distributed Alamouti space-time (ST) codes [2] - when APs 1 and 2 employ the 2×1 Alamouti code to client 1, the resulting diversity gain allows the SNR at client 1 to scale as $|h_{11}|^2 + |h_{21}|^2$, where h_{ij} is the complex chan-

nel gain between AP i and user j . A simpler form of transmit diversity is to transmit the same version of the data stream from multiple transmitters as shown in Fig. 3, wherein client 1 receives a coherent combination of the streams over a composite channel $\tilde{h} = h_{11} + h_{21} + h_{31}$. The power pooled from the multiple transmitters contributes to a combining (SNR) gain on average. This latter form of transmit diversity is similar in principle to broadcast and is often referred to as traditional distributed antenna systems (DAS). Further, unlike ST diversity, DAS does not require the receiver to estimate the individual channels from different transmitters (and associated pilot overhead). This has made it a popular transmit strategy for deployments in stadiums, universities, casinos, hospitals, etc. [1] for both WiFi and cellular signals.

Allowing the data to be accessible from multiple transmit points simultaneously not only provides coverage (during mobility) but also a diversity gain. Further, it does not use CSI (between APs and clients) and hence does not rely on feedback from clients. We refer to such schemes broadly as *diversity* schemes.

Multiplexing: In multiplexing, multiple independent data streams are transmitted concurrently to different users by converting interference into a multiplexing gain through transmitter cooperation - a classic example of which is network (distributed) MIMO [6, 3]. The data streams for different users are shared at each of the transmission points, which are in turn tightly synchronized (at the level of symbol phases). From the PHY layer perspective, this can be realized using a precoding algorithm called zero-forcing beamforming (ZFBF); this applies a precoding matrix (V , computed from channel matrix inverse) to send a linear combination of the data streams through each AP, such that unwanted streams (interference) cancel each other at each client, thereby leaving only the desired stream. A simple scenario is shown in Fig. 2. Network MIMO (netMIMO) can allow the capacity to scale with the number of cooperating transmitters. However, this comes at the cost of leveraging CSI that needs to be fed back from the clients in a timely manner and tight phase synchronization between APs.² We refer to variants of these netMIMO schemes as multiplexing schemes.

Objective: Several works have looked at the practical realization of reuse [8], diversity [5], and multiplexing [3, 6] strategies with a focus on static clients with stable channels. This is understandable as a first step. However, moving forward, given the heterogeneity of user profiles that a system must cater to in both enterprises and outdoor cellular (e.g., small cell LTE, WiMAX) networks, it becomes important to understand which strategies are appropriate for which user profiles and how to intelligently combine them.

3. MAPPING STRATEGIES TO USER PROFILES

²While recent theoretical works [4] are exploring how to effectively leverage outdated CSI, their use in netMIMO is still in its early stages and is hence not considered here.

3.1 User Profiles

Users in enterprise networks can be categorized into one of the following three categories based on their mobility and channel coherence characteristics.

1. *Mobile users with Short Coherence Time:* Coherence time (T_c) varies based on the speed of mobile client which can range from walking speed of 3-4 Kmph ($T_c = 10$ ms) to vehicular speeds of 75 Kmph ($T_c = 1.1$ ms).
2. *Static Users with Short Coherence Time:* Static users can also experience short T_c due to a dynamic environment where objects or other subjects (users) are mobile.
3. *Static Users with Long Coherence Time:* Clients in static environments and in the absence of mobility experience a more stable channel (longer T_c).

3.2 Implementation

In order to evaluate the gains of different transmission strategies for different user profiles, we have implemented the Cooperative (netMIMO and DAS) and the Non-Cooperative reuse (centralized CSMA) schemes on the WARP platform using the WARPLab framework. Implementation of cooperative transmission schemes can be difficult since they require perfect synchronization (w.r.t packet transmission time, sampling clock-rate, center frequency and phase) between the transmitting APs in order to achieve multiplexing gains. In this work, our goal is not to build a large network MIMO or DAS system by overcoming the synchronization challenges between the distributed APs (addressed by other works [6]), but to study the relative performance benefits of these schemes for different user profiles. Since a WARP board supports transmission using multiple radio boards (up to 4), which are tightly synchronized with the internal clock of the main board, these can be directly used for cooperative transmission (netMIMO and DAS) by extending the radio board antennas by 30 ft using LMR400 50 ohm coaxial cable.

Several implementation details (channel estimation, precoding using ZFBF, modulation, receiver decoding, etc.) are not discussed and are given in [9] in the interest of space.

3.3 Experimental Study

Experimental Setup: Our experimental scenario shown in Figs. 1 to 3 is deployed in an indoor lab with three transmitters and three clients. All clients are single antenna WARP boards. To avoid interference all experiments are performed during night and on channel 14 (unused by other devices) in the 2.4 Ghz band. All nodes are connected to a central controller (a PC running the WARPLabs PHY layer signal processing modules) via an Ethernet switch. Reported results are averaged over multiple runs.

Mobility Vs. Transmission Strategy: For mobility we placed the WARP clients on a cart and we moved the cart at walking speeds. We tried to move the cart at the same walking speed on the same path for all mobile experiments. Fig. 4 shows aggregate network rate for each scheme. Aggregate network rate is calculated using achievable Shan-

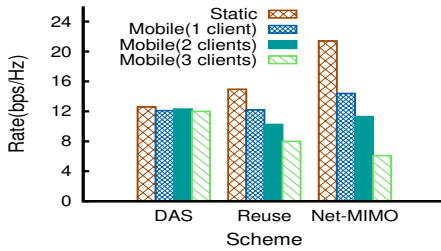


Figure 4: Static vs. Mobile

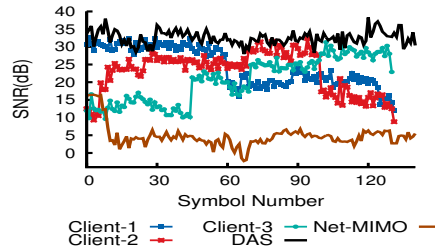


Figure 5: Impact of Mobility

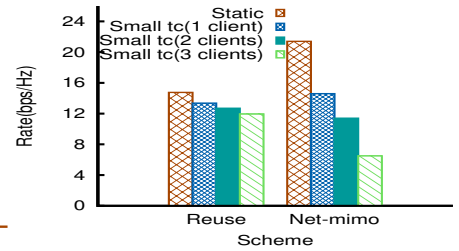


Figure 6: Static with Small Tc

non rate based on received client SINR for each schedules in Figs. 1 to 3. It can be seen that *when all the clients are static, netMIMO is the most appropriate strategy*, outperforming CSMA and DAS by up to 69.8%. netMIMO achieves high network rate by multiplexing three data streams to three clients at the same time. Since the channel coherence time in our static environment is large, all clients can decode their data with high reliability. Further, *since the benefit of reuse outweighs the diversity gain in static environments, even CSMA is better than DAS*. However, as we vary the number of mobile clients in the network, the performance of both netMIMO and CSMA start to degrade; the degradation for netMIMO being especially severe. On the other hand, *owing to coverage and combining gain in SNR from three transmitters (transmitting the same data), DAS is unaffected by client mobility*. Thus, DAS outperforms netMIMO and CSMA by upto 96.7% depending upon the number of mobile clients.

To better understand the behavior of these transmission schemes, we also recorded the received SINR of the symbols transmitted from each transmitter. Fig. 5 shows the SINR of the mobile clients under CSMA, DAS and netMIMO for a single run. For DAS and netMIMO, we only report the SINR values of a single client because other clients also exhibit a similar trend. In the CSMA scheme, client 1 is associated with AP-1, client-2 with AP-2, client-3 with AP-3. During experiments client-1 is moved from left to right while client-3 is moved from right to left. Client-2 is moved in both directions. It can be inferred from Fig. 5 that each client in CSMA, experiences a high SINR only when they are near their respective APs. In DAS, a mobile client experiences a high SNR throughout the experiment due to the coverage and the signal combining effect from the three transmitters, with about 5-6 dB of SINR gain over the highest SINR possible with CSMA (with ideal handoffs). Further, *in addition to link degradation, mobility also impacts the benefits of reuse in CSMA*. In netMIMO, the SINR goes down as soon as the client becomes mobile. Since precoding employs CSI to remove the interference between the concurrent streams, *stale CSI during mobility has a more pronounced impact on netMIMO performance*. Thus, DAS is the best suited scheme for mobile clients. It is of interest to note that while increasing the CSI feedback frequency can potentially reduce the staleness of CSI during mobility, the resulting increase in overhead to keep with up-to-date CSI could become pro-

hibitive (more so with higher # of antennas) and would come at the cost of throughput.

Static Clients with Small Coherence Time: A client can experience fluctuating wireless channel conditions even when its not mobile due to various reasons (such as mobility in environment, multi path etc.). This fluctuation in channel conditions can result in a small channel coherence time even though the client is static. In order to emulate channel variations without moving the clients significantly from their positions, we moved the antenna in the proximity of its original location. Note here that since DAS was the poorest in terms of capacity for static clients, we focus on netMIMO and reuse here. Fig. 6 shows the aggregate network rate achieved by each scheme for clients with channel variations. netMIMO is still highly susceptible to channel variations and the rate degrades in a similar fashion as in mobile scenario. However the corresponding degradation in CSMA is now less pronounced. While link quality is impacted, since the topology does not change, reuse is not impacted and thus, the performance of CSMA is retained. However, stale CSI causes a performance degradation with netMIMO. Hence, interestingly, *a simple reuse scheme such as CSMA serves the best for static clients that experience fluctuating channel conditions (small coherence time)*. Note that since DAS and reuse are both open loop schemes (not using CSI), with the former emphasizing diversity at the expense of reuse, DAS must be cautiously employed for users (eg. mobile) only when reuse is not conducive. While our experiments have been conducted with 3 transmitters, the findings would generalize to larger topologies, where a clustering approach is typically adopted for the realization of cooperative transmission strategies (see sec. 4.1.3).

4. DESIGN OF TRINITY

Two key challenges arise in realizing a practical system that can leverage our inferences: (i) How to categorize users into various profiles? and (ii) How to intelligently combine various strategies to cater to a heterogeneous set of users simultaneously and manage resources effectively between strategies?

4.1 Design Elements

Multiplexing Strategies in Frequency Domain: As mentioned before, the ideal scenario would be to partition the network into disjoint regions, where only one strategy

needs to be applied in each region. This would allow frequency and time resources to be reused by strategies across the network. While such scenarios can occur (e.g., a big conference hall with static users on one end and a cafeteria with mobile users on the other end of a floor), they are not common. In reality, users of different profiles are inter-twined in various regions of the network (e.g., static and mobile users in a cafeteria). Hence, it becomes inevitable to multiplex different strategies either in the time or frequency domain to serve users of different profiles in any given region.

TRINITY employs multiplexing strategy in the frequency domain, which allows it to leverage power pooling benefits from the transmitters that are otherwise not available with time domain multiplexing. In an OFDM system with say N sub-carriers, these sub-carriers would be split between the various strategies in TRINITY (eg., Fig. 9). For example, let N_m sub-carriers be allocated to netMIMO, N_d for DAS and the remaining $N - N_m - N_d$ for reuse. On the downlink (AP→users), when an AP (with fixed transmit power) has users that do not cover all the profiles, then the unused power on the sub-carriers assigned to the unused strategie(s) will be pooled to the sub-carriers assigned to the strategies in operation(results in higher SNR). Such an effect is more pronounced on the uplink (users→AP), where multiple users are served by a given AP simultaneously. Note that when strategies are multiplexed in the time-domain, all sub-carriers are used for a given strategy at a time and hence there is no room for power pooling.

4.1.1 Categorization

We use a combination of CSI as well as sensor hints (eg. accelerometer readings) to help accurately categorize users. Using only one of them (either CSI or sensor hints) is not sufficient to distinguish between the 3 categories.

netMIMO Vs. non netMIMO: Given the difficulty in differentiating user profiles directly from CSI feedback and its fluctuations, TRINITY employs a reactive approach coupled with sensor hints from the users.

When a user joins the network, it begins by aggressively assuming the non-categorized user to be a netMIMO user. Then, based on rate measurements, it reactively determines if the user’s channels allow for netMIMO gains to be leveraged as follows. When channels are measured from the transmitters to users and the precoding matrix is computed for netMIMO, TRINITY keeps track of the estimate of SINR it expects the users to see when netMIMO is executed. In addition, it can also estimate the SNRs the user would receive if the transmitters were to instead operate in DAS and reuse (single transmission per user) modes. The estimated SINR can easily be calculated using the proposed technique in [10]. When netMIMO is executed, the resulting SINR or rate is then measured and compared against the estimated value. If there is a significant difference between the estimated and observed rates for a user and is less than the rate estimated for DAS, then the user is removed from the netMIMO category. However, if there is a degradation in the

observed netMIMO rate but it is still higher compared to that resulting from DAS, then it makes sense to retain the user in the netMIMO category.

To understand the validity of our claim, we simulate a scenario with different channel coherence times to create mobility (0 to 75 Kmph). Instead of transmitting the symbols over the air we pass them through a flat AWGN channel, which remains constant over the coherence time (channel feedback rate is 100ms) and then changes independently to a new realization. It is seen from Fig. 7 that netMIMO rate drastically degrades for a client when it changes its state from being static to mobile at walking speed. Unlike DAS or CSMA, the deviation between the estimated and measured rates is very large for netMIMO (due to the reliance on CSI). This gives us confidence in categorizing netMIMO users based on rate discrepancies.

DAS Vs. Reuse (CSMA): Among the users in the non-netMIMO category (with small coherence times), to distinguish between ones that require DAS from those that require reuse, TRINITY employs sensor hints (similar to [7]) in the form of accelerometer readings from the users. Based on the degree of mobility predicted by the sensor hints and the density of transmitter deployment, TRINITY can estimate the potential frequency of handovers for the user, and hence determine the appropriateness of DAS (mobile user) or reuse (very low mobility or static user with environment dynamics) strategies for the user. Note that all devices may not be capable of providing sensor hints. However, users that are mobile would invariably access data through their smartphones, providing sensor hints from which is not an issue. Hence, for users in the non-netMIMO category that are unable to provide sensor hints, we can categorize them into the reuse profile with high probability. Note that sensor hints can also be used as a pre-processing step to filter out the highly mobile (and hence non-netMIMO) users.

Re-categorization: Note that a user’s profile can change from time to time. Hence, to keep track of user dynamics, TRINITY periodically moves a user in the non-netMIMO category to netMIMO category and re-categorizes it using the above procedure. However, for users in the netMIMO category (for whom CSI is available), as and when they see performance degradation, they can be immediately moved to the appropriate category based on the above procedure.

4.1.2 Resource Management

Once the users are categorized, the traffic load for each of the strategies can be determined based on the traffic carried by users in the respective category. The different user profiles may be weighted (based on priority or fairness), and the allocation of number of sub-carriers to each of the strategies can be made proportional to their weighted traffic load. Among the sub-carriers allocated to a strategy, users in the respective profile can be scheduled based on any fairness model (eg. proportional fairness).

4.1.3 Strategy Optimization

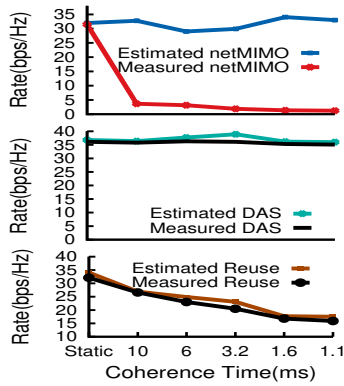


Figure 7: Estimated vs. Measured

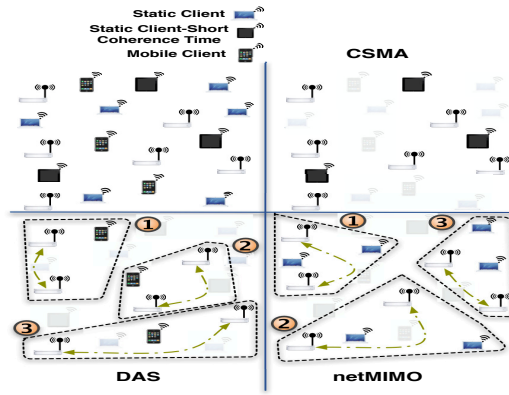


Figure 8: Optimizing Strategies

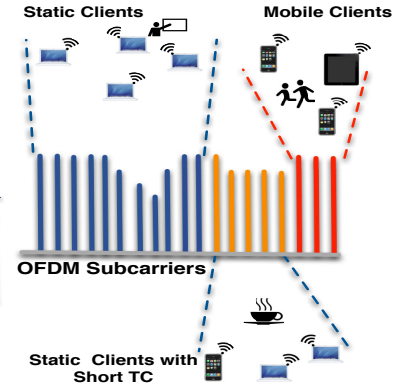


Figure 9: Strategy Multiplexing

Once sub-carriers are assigned to the different strategies, TRINITY optimizes the execution of each of the strategies to their respective clients.

Optimizing netMIMO: For netMIMO, the ideal operation would be to execute one large netMIMO between all the transmitters in the network and the users in the associated category. However given the complexity, this is not desirable, especially for large networks, where the central controller manages several tens of transmitters. To strike a balance between performance and complexity, TRINITY decomposes the set of transmitters into smaller, contiguous clusters of transmitters (Fig. 8), wherein netMIMO is executed only within each cluster and interference between clusters is avoided either in the time or frequency domain.

Optimizing DAS: Unlike netMIMO, the tradeoff seen in the DAS mode is that between coverage and capacity. Similar to netMIMO optimization, TRINITY employs a clustered approach to strike a fine balance between coverage and capacity (Fig. 8), wherein DAS is employed only within smaller, contiguous clusters of transmitters. This allows sub-carriers to be reused across DAS clusters (e.g., cluster 1 and 3 can operate simultaneously), subject to interference avoidance between clusters.

Optimizing Reuse: Optimizing the reuse strategy amounts to maximizing reuse in conventional wireless (eg., CSMA) networks. Several solutions have been proposed (eg., [8]) in this context and can be adopted for the reuse strategy in TRINITY.

Additional details on optimizing different strategies are given in [9] in the interest of space.

5. DISCUSSIONS AND CONCLUSIONS

Multiplexing users in the frequency domain is common to OFDMA systems and can hence be implemented efficiently. Employing time domain for multiplexing would reduce the complexity further, however at the expense of power pooling benefits. The bulk of the overhead in TRINITY arises from CSI feedback, which is the price to pay for any closed-loop MIMO gains. However, note that, with the help of sensor hints, we restrict such fine grained CSI feedback only for

appropriate users (and not all users) for whom netMIMO can be enabled, thereby keeping the overhead reasonable.

While most of our discussions have been *w.r.t.* downlink, our solutions would apply to uplink as well (clients→APs). However, unlike downlink, uplink OFDMA requires sub-carrier level synchronization between clients, which is challenging but do-able (eg. LTE, WiMAX). More importantly, TRINITY is applicable to upcoming, outdoor small cell (LTE, WiMAX) networks as well (not just enterprise networks). Indeed, the problem of mobility is exacerbated in outdoor small-cell networks, wherein the potential benefits of TRINITY would be even more pronounced.

To summarize, we envision TRINITY to identify and optimize transmission strategies that can cater to users of various profiles effectively. We believe such an approach is critical to improving user quality of experience in next generation wireless networks, where user (device) heterogeneity will be the norm.

6. REFERENCES

- [1] DAS in Action : Atlanta. <http://bit.ly/ACRt4V>.
- [2] ALAMOUTI, S. M. A Simple Transmit Diversity Technique for Wireless Communications. In *IEEE Journal on Select Areas in Communications* (1998).
- [3] LIN, K., GOLLAKOTA, S., AND KATABI, D. Random access heterogeneous mimo networks. In *ACM SIGCOMM* (2011).
- [4] MADDAH-ALI, M., AND TSE, D. Completely stale transmitter channel state information is still very useful. In *Allerton* (2010).
- [5] RAHUL, H., HASSANIEH, H., AND KATABI, D. Sourcesync: a distributed wireless architecture for exploiting sender diversity. In *ACM SIGCOMM* (2010).
- [6] RAHUL, H., KUMAR, S. S., AND KATABI, D. MegaMIMO: Scaling Wireless Capacity with User Demand. In *ACM SIGCOMM* (2012).
- [7] RAVINDRANATH, L. S., NEWPORT, C., BALAKRISHNAN, H., AND MADDEN, S. Improving Wireless Network Performance Using Sensor Hints. In *USENIX NSDI* (2011).
- [8] SHRIVASTAVA, V., AHMED, N., RAYANCHU, S., BANERJEE, S., KESHAV, S., PAPAGIANNAKI, K., AND MISHRA, A. Centaur: realizing the full potential of centralized wlns through a hybrid data path. In *ACM MobiCom* (2009).
- [9] SINGH, S., SUNDARESAN, K., KHOJASTEPOUR, M., RANGARAJAN, S., AND KRISHNAMURTHY, S. One strategy does not serve all: Tailoring wireless transmission strategies to user profiles. In *NEC Labs America Technical Report* (2012).
- [10] TARIGHAT, A., SADEK, M., AND SAYED, A. A multi user beamforming scheme for downlink mimo channels based on maximizing signal-to-leakage ratios. In *IEEE ICASSP* (2005).