ProBeam: A Practical Multicell Beamforming System for OFDMA Small-cell Networks

Jongwon Yoon*
University of Wisconsin, Madison, WI
yoonj@cs.wisc.edu

Karthikeyan Sundaresan
NEC Labs America Inc.
Princeton, NJ
karthiks@nec-labs.com

Mohammad (Amir) Khojastepour
NEC Labs America Inc.
Princeton, NJ
amir@nec-labs.com

Sampath Rangarajan
NEC Labs America Inc.
Princeton, NJ
sampath@nec-labs.com

Suman Banerjee
University of Wisconsin, Madison, WI
suman@cs.wisc.edu

ABSTRACT

Small cells form a critical component of next generation cellular networks, where spatial reuse is the key to higher spectral efficiencies. Interference management in the spatial domain through beamforming allows for increased reuse without having to sacrifice resources in the time or frequency domain. Existing beamforming techniques for spatial reuse, being coupled with client scheduling, face a key limitation in practical realization, especially with OFDMA small cells. In this context, we argue that for a practical spatial reuse system with beamforming, it is important to decouple beamforming from client scheduling. Further, we show that jointly addressing client association with beamforming is critical to maximizing the reuse potential of beamforming.

Towards our goal, we propose *ProBeam* – a practical multi-cell beamforming system for reuse in small cell networks. ProBeam incorporates two key components - a low complexity, highly accurate SINR estimation module that helps determine interference dependencies for beamforming between small cells; and an efficient, low complexity joint client association and beam selection algorithm for the small cells that accounts for scheduling at the small cells without being coupled with it. We have prototyped ProBeam on a WiMAX-based network of four small cells. Our evaluations reveal the accuracy of our SINR estimation module to be within 1 dB, and the reuse gains from joint client association and beamforming to be as high as 115% over baseline approaches.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless Communication

*Jongwon Yoon interned in the Mobile Communications and Networking department at NEC Laboratories America Inc, Princeton during this work.

MobiHoc'13, July 29–August 1, 2013, Bangalore, India.

General Terms

Algorithms, Design, Experimentation, Performance

Keywords

Beamforming, OFDMA, WiMAX, Small-cell

1. INTRODUCTION

The proliferation of smartphones and tablet devices has made it necessary for mobile operators to consider new technologies that provide increased network capacity. Small cells (micro and pico cells) provide a promising solution to address this need and are already being deployed for 3G networks, with future rollouts of 4G small cells [1]. With reduced cell sizes and dense deployments, small cells are geared for increased spatial reuse of spectral resources – a valuable and scarce commodity in next generation wireless networks (WiMAX, LTE, LTE-advanced, etc.).

Given the dense deployment of small cells, interference plays a key limiting factor in harnessing their potential. While the sheer scale limits planned deployment of small cells (similar to WiFi), handling interference is a very different problem in small cells compared to WiFi. This can be attributed to their synchronous access mechanism (borrowed from macrocells), coupled with OFDMA (orthogonal frequency division multiplex access) transmissions, wherein multiple users are served in the same frame. Earlier works on interference management in small cell networks [2, 3] employed interference avoidance in the time or frequency domain by allocating orthogonal resources to interfering small cells. In this work, we aim to avoid such sacrifices of spectral resources by exploring interference management for small cells in the spatial domain through beamforming antennas.

Employing beamforming or directional antennas for spatial reuse in a multi-cell set-up has been considered in the context of WiFi [4, 5]. However, such approaches face a key limitation when it comes to practical realization in that a single client is assumed for each AP when computing interference conflicts and determining the spatial reuse schedule. When the client scheduled for an AP changes, the interference conflicts change, requiring a re-computation of the schedule, potentially at the granularity of every packet. This makes it hard to realize such solutions for practically sized WiFi networks and *more so for small-cell networks, where multiple clients are scheduled in each OFDMA frame*. Hence, the goal of this work is to leverage beamforming for spatial reuse across small cells but at the same time decouple it from per-frame scheduling at the small

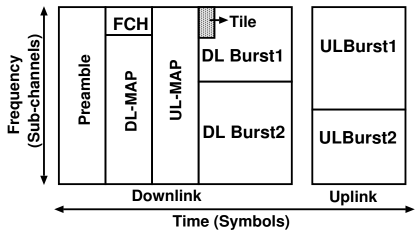

Figure 1: WiMAX frame structure.

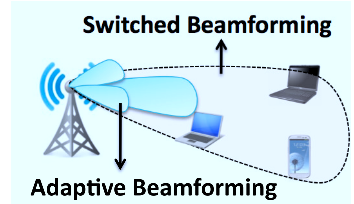

Figure 2: Illustration of beamforming.

cell base station (BS), thereby allowing for beam selections to be computed only at the granularity of seconds (hundreds of frames).

Executing beam selection at coarser time scales compared to client scheduling allows for tangible spatial reuse benefits across cells. However, the beam chosen for a small cell must now deliver good transmission rates to *all* the users that are associated (and hence can be scheduled) with the small cell in order to realize the throughput gains from spatial reuse. Hence, we argue that to realize practical and efficient spatial reuse with small cells, it is important to not just decouple beam selection from scheduling but also integrate beam selection with client association. Towards this goal, we propose *ProBeam* – a practical system that enables joint multi-cell beamforming and client association for increased spatial reuse in small cell networks.

ProBeam incorporates two key components: (i) a *SINR estimation module* – this captures the interference dependencies between small cells in the presence of beamforming. Note that accurate SINR estimation would require measurement *w.r.t* all possible combination of beam choices at small cells, resulting in $O(k^n)$ measurements, where k and n are the number of beam choices and small cells respectively. ProBeam's estimation module indirectly computes SINR from SNR measurements, thereby resulting in only linear number of measurements ($O(kn)$) with an estimation error less than 1 dB for 95% confidence and a maximum error of 1.65 dB. (ii) a *joint beam selection and client association module* – given the hardness of beam selection and client association problems in isolation, their joint problem is significantly challenging to address optimally. ProBeam employs an efficient yet greedy $\frac{1}{2}$-approximation algorithm for client association as a building block to converge to an efficient spatial reuse solution with both beam selection for small cells along with their client associations.

We have implemented ProBeam on a four cell WiMAX-based small cell network. Our experimental evaluations reveal that ProBeam is within 90% of the optimal solution and provides close to 50% throughput gains by addressing the joint problem of client association along with beam selection compared to existing approaches that address only the latter.

Our contributions in this work are multi-fold.

- We propose a low, linear complexity SINR estimation scheme with an error less than 1 dB to generate the interference dependencies needed for computing spatial reuse configurations.
- We establish the hardness of the joint beam selection and client association problem and propose a practical, yet efficient algorithm to address the same.
- We demonstrate the practicality and showcase the benefits of ProBeam by prototyping and evaluating it on a WiMAX-based network of four small cells.

The rest of the paper is organized as follows. Section 2 provides background on OFDMA systems and related work. We motivate the need to couple client association with multi-cell beamforming in Section 3. We describe the algorithm in Section 4 and evaluate the performance of ProBeam using WiMAX testbed in Section 5. Section 6 concludes the paper.

2. BACKGROUND AND RELATED WORK

2.1 WiMAX Preliminaries

OFDMA small cells: Next generation small cell networks for LTE and WiMAX borrow their access mechanism from their macrocell counterparts and are based on OFDMA. Further, they operate on licensed spectrum and follow a synchronous access mechanism (unlike WiFi), wherein frames are transmitted periodically at fixed time intervals (1 ms in LTE, 5 ms in WiMAX). Each OFDMA frame is a two-dimensional template (time and frequency slots) that carries data to *multiple* clients – another key difference compared to WiFi. Transmissions between downlink (DL, BS to client) and uplink (UL, client to BS) are separated either in frequency (FDD) or in time (TDD). Figure 1 shows an example of a WiMAX TDD frame, the underlying structure of which is common to LTE as well. Every frame carries a control and a data part, where the control part (e.g., DL and UL MAPs) provides information to the clients regarding where to pick (place) their respective downlink (uplink) data from the frame and what parameters (MCS - modulation and coding scheme) to use for decoding (encoding) the downlink (uplink) data. Clients use the uplink frame to report their instantaneous CSI (channel state information) to the BS, which in turn is used for diversity scheduling at the BS.

Given the dense deployment of small cells, resource and interference management among small cells happens at the cluster (tens of small cells) granularity, wherein a central entity (SON: self organizing network server [6]) or one of the small cells in the cluster performs centralized resource management for the cells in the cluster and coordination is achieved with the help of a high speed backhaul. While clients use the preamble and control part of the frame to synchronize to the BS, the small cell BSs themselves can synchronize to the macrocell with the help of the SON server or with a GPS antenna module.

Beamforming: Beamforming is one of the core features in next generation networks that is adopted to improve SNR at the intended receivers while decreasing interference at unintended receivers. A beamforming system typically uses multiple antenna elements in an array to form various beam patterns. Beam patterns reinforce transmission energy in desired directions by weighting the signal from the antenna array in both magnitude and phase. Beamforming can be either switched (directional) or adaptive. In switched, a pre-determined set of directional beam patterns covering the azimuth are stored and chosen based on coarse feedback (SNR or RSSI) from the client. In adaptive, fine-grained feedback of channel estimation from the client is used to adapt the beam pattern on the fly to reinforce multipath components and maximize the SINR at the client. By adapting to the instantaneous multipath channel, adaptive provides higher gain (at the cost of increased feedback) compared to switched. However, at the same time, it is more sensitive to channel fluctuations and requires timely feedback to track the channel state - a limiting constraint especially during mobility and in multi-cell resource management.

Both switched and adaptive beamforming co-exist in a complementary manner in cellular systems. Macrocells are sectored in operation (e.g., three 120° or six 60° directional beams), while adaptive beamforming is enabled to clients within each of the sectors separately. Unlike macrocells, where interference is restricted to cell-edges, thereby allowing for all sectors to operate in tandem, interference is a more pervasive phenomenon in small cells [7]. This requires small cells to select a single sector (switched beam) for operation in a frame (adaptable across frames) so as to avoid interference and maximize reuse among small cells in a dense deployment. Note that adaptive beamforming can still be enabled to clients within the sector of operation at each small cell (see Figure 2 for illustration).

2.2 Related Work

Interference has been shown to be a key performance limiting factor for small cells [7]. This necessitates interference mitigation solutions that incorporate dynamic resource partitioning strategies. There have been studies [8, 9] in this direction but are restricted to theory with several simplifying assumptions that restrict their scope and deployment. Recently [2] and [3] propose centralized and distributed resource management schemes respectively for interference mitigation and demonstrate their efficacy in practice. These solutions allocate orthogonal resources to interfering small cells to avoid interference while reusing resources for the clients that do not incur interference. However, such resource isolation either in time or frequency comes at the cost of sacrificing resources, which in turn can be avoided by addressing interference in the spatial domain through beamforming.

In the space of beamforming, [4, 5] propose to increase the capacity of WLANs through spatial reuse by considering directional antennas only at the APs or at both APs and clients. However, client association is assumed and conflicts and reuse schedule are computed *w.r.t* a single client at each AP. This limits the practical applicability of such solutions (especially for OFDMA systems) since conflicts and reuse schedules have to be recomputed (potentially every packet) every time the client scheduled with any of the AP changes. Several theoretical works [10, 11] have looked at adaptive beamforming in a multi-cell context. However, idealized settings are assumed that require fine grained CSI from all transmitters to all clients be made available to the reuse algorithm at every frame interval. Given the practical feasibility (or lack thereof) of such approaches, experimental works [12] have appropriately focused on adaptive beamforming for SNR improvements within a single cell. Further, none of these works address client association jointly with beamforming.

The focus of our work is to design a *practical* multi-cell spatial reuse system that, decouples client scheduling from beamforming, employs switched beamforming for interference management between small cells, and jointly addresses client association to increase the potential of spatial reuse from beamforming. Being complementary, adaptive beamforming can still be leveraged for SNR improvement within each small cell (although not considered in this work).

3. MOTIVATION

We now motivate the need to couple client association with multi-cell beamforming in order to maximize the benefits of spatial reuse. We present results from an experimental WiMAX-based network of four small cells, each equipped with an eight element phased array antenna (details in Section 5) to substantiate our claims.

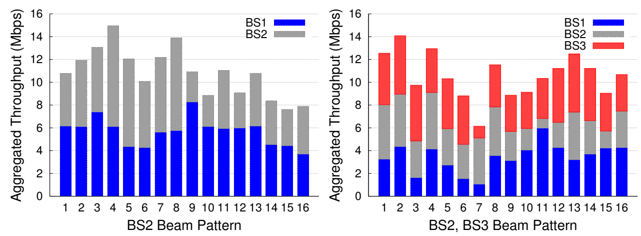

(a) Two cell network (b) Three cell network

Figure 3: Motivation for coordinated beam selection.

3.1 Need for Coordinated Beamforming

Beamforming in a multi-cell context has two benefits: (i) increase link capacity through improved SNR, and (ii) increase network capacity through reduced interference (higher SINR) and hence higher spatial reuse. The beam choice of one cell impacts the interference seen by the clients of another cell, thereby requiring a coordinated approach to beam selection across small cells for maximum reuse benefits. However, given the simplicity of un-coordinated, per-cell beamforming (focusing only on SNR), it is important to understand the benefits from coordination and hence the need for it.

We construct a topology with two small cells, each with one scheduled client. First BS1 cycles through all its sixteen beam patterns to determine the one yielding the best rate to its client (C1) in isolation. BS1 is then fixed to use its best beam to C1. Now, in the presence of BS1, BS2 is made to transfer data to its client (C2) on each of its 16 patterns sequentially. We plot the throughput observed at C1 (blue bars) and C2 (grey bars) as a function of the beam pattern used by BS2 in Figure 3(a). Two observations can be made: (i) The interference projected by BS2 on C1 depends tightly on the beam chosen by BS2. C1 achieves its highest throughput (8.3 Mbps) when BS2 employs its 9th pattern and its lowest throughput (3.7 Mbps) when BS2 employs its 16th pattern. (ii) The beam maximizing the throughput of one cell does not necessarily maximize the multi-cell network throughput. While the 9th beam pattern maximizes C1's throughput, it is the 4th pattern that maximizes the aggregate network throughput. A similar behavior is also evident in the three cell experiment presented in Figure 3(b), where the pattern (11th) maximizing throughput for C1 differs from the one (2nd) maximizing the aggregate network throughput. The throughput gain of employing the 2nd pattern over the 11th one is almost 40%.

Thus, *a well-coordinated beamforming algorithm across the small cells is indeed important to maximize the aggregate network throughput*.

3.2 Need for Joint Client Association

Client association has been traditionally employed to load balance clients between multiple cells so as to effectively utilize the capacity of each cell and network as a whole. However, in the context of multi-cell beamforming, client association has a bigger role to play. Note that, unlike in WiFi systems, where a single client is served by a cell at a time, OFDMA systems multiplex multiple clients in the same frame (diversity scheduling). This requires that

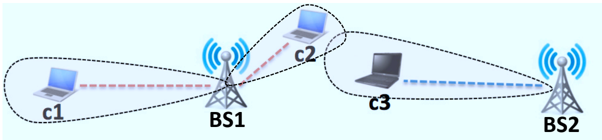

(a) SNR based association (C2 is associated with BS1)

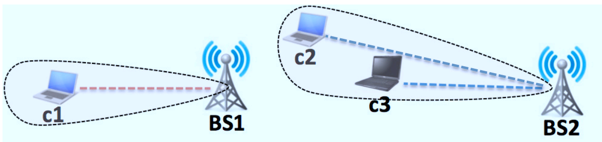

(b) Flexible association (C2 is associated with BS2)

Figure 4: Illustration for flexible client association.

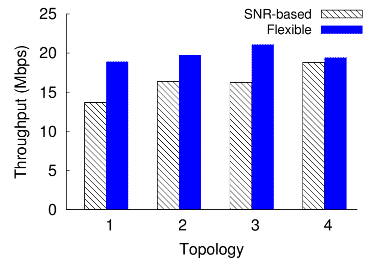

Figure 5: Joint association increases throughput by 40% comparing to decoupled (SNR based) association.

the beam selected for the small cell cater effectively to all its associated and scheduled clients. Further, since the beam choice for a cell impacts the interference and hence performance seen by other cells, this naturally results in client association being closely coupled with multi-cell beamforming.

To see this, consider the following illustration in Figure 4. In conventional association, where SNR is used as a metric for client association, clients C1 and C2 will be associated to BS1, while C3 will be associated to BS2 based on (high) SNR and completely decoupled from beamforming. BSs will then determine the best beams to communicate with their respective clients. Let $b1$ and $b2$ be the only beams on which C2 and C3 can receive good signal strength from their respective BSs. Now, when BS1 is employing beam $b1$ to communicate with C2, this will receive interference from the beam $b2$ used by BS2 to communicate with its client C3. By fixing the client association, depending on the location of associated clients, the ability of beamforming to effectively suppress interference between cells is potentially limited. In contrast, by allowing for flexible association (Figure 4(b)), C2 can be associated with BS2 even though it has a lower SNR to BS2. This would allow BS2 to schedule C2 and C3 jointly on a beam that suffers no interference from that employed by BS1, thereby resulting in a potentially higher SINR for all clients.

To quantify the benefits of coupling client association with beam selection for small cells, we conduct the following experiment with two small cells and three clients, and generate multiple topologies by varying the client locations. We consider two association strategies: *decoupled association*, where the best coordinated beam (for maximum aggregate throughput) for each small cell is selected after client association is done based on SNR; *joint association*, where the client association yielding the highest aggregate throughput is computed among all beam combinations between the two cells. The aggregate throughput results between these two strategies in Figure 5 indicates that joint association can yield gains as high as 40%, with an average gain of about 25%.

This in turn motivates the *need to jointly address client association with beam selection for small cells, whereby client association can be effectively used to maximize the spatial reuse potential of beamforming*.

4. DESIGN

System overview: Small cell networks can be deployed for enterprises as well as outdoors. A central controller (separate entity or one of the small cells) is designated to perform resource and interference management for a cluster (tens) of small cells jointly with a high speed backhaul available for information exchange between them. We expect ProBeam to reside in this central controller (CC). Note that while our primary focus is small cell networks, our system is equally applicable to WiFi networks as well.

ProBeam's spatial reuse solution operates in epochs, which spans several seconds (hundreds of frames). In each epoch, the sequence of operations are as follows. (i) *Interference estimation for beamforming*: The clients measure the average SNR on each of the beams from each of the BSs and forward it to the CC, which then infers their corresponding SINR for various beam combinations at the small cell BSs (details in subsection 4.1). (ii) *Joint beam selection and client association*: Based on the interference information collected, the CC runs its spatial reuse algorithm (for a desired objective) to determine the beam choice for each of the small cells as well as the clients that are associated with it for that epoch (details in subsection 4.2). (iii) Scheduling: Once each small cell BS receives its beam choice and client set, it begins scheduling its clients locally using its own scheduler (proportional fair, max-min fair, etc.) for each frame in the epoch, while applying the beam selected to the frame transmissions (details in subsection 4.3).

4.1 Interference Estimation for Beamforming

Estimating the interference at clients accurately is critical for the efficient operation of ProBeam.

Reducing complexity: Measuring the SINR directly at the clients for various beam configurations (interference) used by small cells is the most accurate approach. However, this would entail that each small cell cycle through each beam pattern, while keeping the beam patterns at other cells fixed and measuring the resulting SINR at all clients. This would however result in a total of $O(k^n)$ measurements, where k in the number of beam patterns and n is the number of small cells. ProBeam measures only the client SNR from each of the small cells in isolation for the various beam choices and then uses this information to estimate the projected client SINR for a given beam configuration at the small cells. By allowing the small cells to operate in isolation during measurements, this significantly reduces the SINR estimation complexity to $O(kn)$. The key question remaining is the accuracy or lack thereof of such an estimation procedure.

Note that SINR can be expressed as $SINR_{ij} = \frac{SNR_{ij}}{\sum_{k \neq i} INR_{kj}+1}$, where SINR at client j from BS i is related to its SNR and net interference to noise ratio from other BSs ($INR_j = \sum_{k \neq i} INR_{kj}$). Small cells being interference limited, $INR + 1 \approx INR$. In the logarithmic (dB) domain, the relation can be expressed as SINR (dB) = SNR (dB) - INR (dB). Hence, in principle, the SINR at a client can be estimated from its SNR from the desired BS and its aggregate INR from all interfering BSs. For this to be possible, one needs to estimate each INR_{kj}, which can potentially be approximated as the client SNR when associated with the interfering BS in

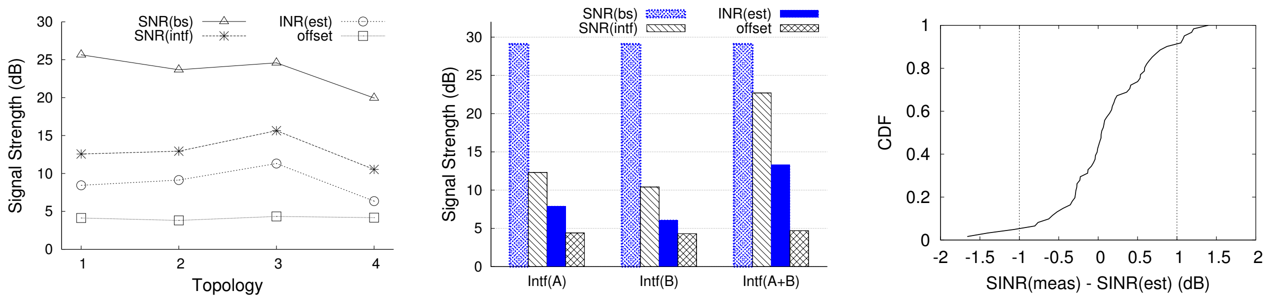

(a) Interference estimation (b) Estimation with 2 interferers (c) CDF of estimation error

Figure 6: Accurate estimation of SINR from individual SNR estimates.

isolation (i.e., SNR_{kj}). However, in reality, the accuracy of such an estimation may depend on multiple factors such as quantization, offsets, estimation error, etc.

To verify this approximation, we conducted the following experiment with two small cell BSs, whose beam choices are such that they interfere with the client under consideration. The results are presented in Figure 6(a). First, the client measures the signal strength from the associated BS in the absence (SNR_{BS}) and presence ($SINR_{BS}$) of interference respectively, from which we estimate $INR_{est} = SNR_{BS} - SINR_{BS}$. Then, the client records the signal strength SNR_{intf} after associating with the interfering BS in isolation. Comparing SNR_{intf} with INR_{est} in Figure 6(a), we see that there is a consistent 4 dB offset between the estimated interference and its corresponding signal strength and this remains fixed regardless of the topology and client SNR considered (SNR_{BS} range 19 to 26 dB). We attribute this constant 4 dB difference to the inherent offset β introduced (during client feedback) by the MAC and its quantization of the signal strength value reported from the PHY layer. β being platform dependent, can be calibrated by the client and fed back to the Central Controller for its appropriate estimation of INR. Further, note that when SINR is directly measured, there is only one feedback value from the client. However, when SINR is estimated from SNR and multiple INRs, then each of the SNR feedback (corresponding to INR) introduces an offset that needs to be compensated. When appropriately compensated, the resulting estimation reduces to $SINR_{ij}$ (dB) $= SNR_{ij}$ (dB) $- 10\log_{10}(\sum_{k\neq i} SNR_{kj}) + \beta$. Note that since interference is aggregated in absolute units, the offset for the aggregate interference remains to be β in dB. This is observed in Figure 6(b), where the offset in the presence of one (A or B) and two (A+B) small cell interferers remains to be the same 4 dB. Thus, with the help of isolated measurements from the small cells, it is indeed possible to estimate SINR, thereby resulting in a linear (in n) complexity of only $O(kn)$.

SINR estimation procedure: ProBeam initiates a measurement phase at the beginning of each epoch, where it operates each small cell BS in the cluster one after another in isolation. When activated, BS i applies its k beam patterns sequentially, each lasting ten frames. All the clients measure the average received SNR from BS i corresponding to beam pattern k. A client j forwards SNR_{ijk}, i.e., measured SNR from BS i with beam pattern k to the CC in ProBeam through its current associated BS. In WiMAX and LTE, clients automatically send Channel State Information (CSI) to BS periodically via dedicated uplink channel resources in every frame. We use such standard feature for obtaining our desired SNR measurements. Once ProBeam gathers SNR measurements from all the clients, then any desired SINR (in dB) for a given beam configuration ($\pi = \{\pi(1)\},\ \forall i$, beam choices for small cells) can be estimated as,

$$SINR_{ij\pi}(dB) = SNR_{ij\pi(i)}(dB) - 10\log_{10}\Big(\sum_{k\neq i} SNR_{kj\pi(k)}\Big) + \beta(dB) \quad (1)$$

Note that SNR measurements can be done within $kn \times 10$ frames. For reasonable values of k (say 10 beams) and n (say 10 cells in a cluster), this would amount to 1 sec in LTE (for 1 ms frames). Also actual data is transmitted during the measurement phase, therefore we do not waste resources for SNR measurements. However, reuse cannot be leveraged, whose overhead (reuse loss) can be amortized as long as the epoch duration is several seconds.

Validation: To validate our estimation procedure, we conduct the following experiments with three small cell BSs and a single client. First, the client measures the SNRs from all three BSs for a given beam configuration in isolation and records them. Then, we make the client associate with one of the BS and measure the SINR in the presence of the other two BSs projecting interference. The beam configuration is chosen so as to project interference to the client under consideration. We repeat the above experiment by changing the beam configuration as well as the topology (i.e., client locations) to obtain confidence in results. Measurements are taken at different client locations to generate plurality of interference scenarios and to also emulate different clients (varying BS deployment is considered in Section 5). We obtain over 100 sets of measurements and present the CDF of the SINR estimation error ($SINR_{meas} - SINR_{est}$) in Figure 6(c). As we can see, 95% of our SINR estimates have less than 1 dB error ($\leq$ 5%), with the highest estimation error being only about 1.65 dB. *Our results clearly indicate the high accuracy of our SINR estimation method, thereby avoiding the complexity of obtaining measurements from all possible combinations of beam patterns at small cells.*

4.2 Joint Client Association and Beam Selection (CABS)

Similar to other resource management problems, we can formulate our problem as a utility maximization problem in every epoch.

$$\text{Maximize} \sum_{j\in\mathcal{K}} U(t_j)$$

where t_j represents the average throughput received by client j in the epoch and $U()$ is a function to capture the corresponding utility. Note that the choice of the utility function determines the fairness policy in the system. We assume utility functions to be concave and non-decreasing. This captures proportional fairness (defined by using the utility function $U(t_j) = \log(t_j)$) that is popular in the standards (WiMAX, LTE). While we need to decouple the time scales of operation for CABS from scheduling, it must be noted that the eventual objective is related to throughput and hence dependent on scheduling. Hence, to allow the decoupling, throughput needs to be modeled as the average throughput received by the client over the epoch for a given scheduling policy. Our problem can be formulated as,

$$(\pi^*, \mathbf{x}^*) = \arg\max_{\pi, \mathbf{x}} \sum_{j \in \mathcal{K}} \sum_{i \in \mathcal{S}} x_{ji} U(t_{ji}^{\pi}) \text{ s.t. } \sum_{i \in \mathcal{S}} x_{ji} \leq 1, \ \forall j \in \mathcal{K} \tag{2}$$

where $\mathcal{K}$ and $\mathcal{S}$ represents the set of clients and small cell BSs respectively. Further, $\pi = \{\pi_i, \ \forall i\}$ denotes the beam selection vector for all BSs, while $\mathbf{x} = \{x_{ji}, \ \forall j, i\}$ denotes the association vector for all clients ($x_{ji} = 1$ if client j is associated with BS i and 0 otherwise). t_{ji}^{π} indicates the client j's average throughput when associated with BS i under beam configuration π and depends on the SINR ($SINR_{ij\pi}$) seen by the client from BS i in the presence of interference from other BSs under the beam configuration π (see Eq. (1)).

4.2.1 A note on fairness

While fairness (starvation) among clients is typically achieved (avoided) over a longer time period, instantaneous per-frame decisions may favor clients with good channel conditions (e.g., proportional fairness). In the case of CABS, decisions are made at the granularity of epochs. Hence, if fairness is ensured over much longer time scales (>> epoch), then several clients could be subject to starvation in an epoch (several seconds). This would increase the jitter perceived by such clients – a factor critical for real-time media and is hence not desired. Thus, it is more appropriate to ensure fairness within each epoch. This would allow all clients to be scheduled in every epoch. On the other hand, since beam selection decisions are fixed for the entire epoch, accommodating all clients could potentially limit the amount of reuse that can be leveraged in the epoch. Hence, to strike a balance between throughput performance (reuse) and fairness, an alternative is to restrict the utility functions to be non-negative in addition to concave and non-decreasing. This would account for fairness, while at the same time allowing for a small number of clients to be removed from scheduling in an epoch. By weighting the client utility functions inversely proportional to their throughput received (T_j) thus far, one can avoid starvation for all clients across epochs.

In the case of proportional fairness, we can modify the utility function as $U(t_j) = w_j \log(t_j)$; if $t_j > 0$ and 0 otherwise, where $w_j \propto \frac{1}{T_j}$. Further, T_j at current epoch e is updated through an exponentially weighted moving average as $T_j(e) = (1 - \frac{1}{\alpha})T_j(e-1) + (\frac{1}{\alpha})t_j(e)$, where α is the filtering coefficient. Let r_{ji}^{π} be the average transmission rates (MCS) seen by client j in a slot when associated with BS i under beam configuration π, and N be the total number of time-frequency slots in an OFDMA frame with M frames per epoch. Then, under proportional fairness, it can be easily shown that the number of slots are allocated among all the scheduled clients in the proportion of their weights (equal when $w_j = 1, \forall j$). This would in turn result in an average client throughput of $t_{ji}^{\pi} = \frac{NMw_j r_{ji}^{\pi}}{\sum_{k \in \mathcal{K}} x_{ki} w_k}$.

4.2.2 Hardness

For a given client association, the problem of beam selection is itself NP-hard [4, 5]. Hence, it comes as no surprise that our joint CABS problem is NP-hard as well. From the perspective of designing algorithms, it helps to understand if beam selection is the only source of hardness or does client association also contribute to the hardness. In this regard, we have the following result.

THEOREM 1. *For a given beam selection, the CABS problem remains to be NP-hard.*

In the interest of space, we defer the proof to [13].

Algorithm 1 CABS Algorithm

```
1:  INPUT: average SNR ρ^b_ji, ∀i ∈ S, j ∈ K, b ∈ B
2:  OUTPUT: Beam selection π(i) and client association A_i, ∀i ∈ S
3:  Initialization of beam choices, i.e., π(i), ∀i
4:  for i ∈ [1 : |S|], b ∈ [|B|] do
5:      L = ∅, u_ib = 0
6:      while 1 do
7:          j* = arg max_{j∈K\L} Σ_{k∈L∪j} U(t^b_ki) − u_ib
8:          if j* = ∅ then break
9:          L ← L ∪ j*; u_ib = Σ_{k∈L} U(t^b_ki)
10:     end while
11: end for
12: π(i) = arg max_b u_ib, ∀i
13:
14: for i ∈ [1 : |S|] do
15:     for b ∈ [1 : |B|] do
16:         % Solve client association by varying only one beam element at a time
17:         π(i) = b, A_i = ∅, ∀i
18:         (i*, j*) = arg max_{(i,j) s.t. j∉∪_i A_i} {Σ_{k∈A_i∪j} U(t^π_ki) − Σ_{k∈A_i} U(t^π_ki)}
19:         A_i* ← A_i* ∪ j*; u^π_ib = Σ_i Σ_{j∈A_i} U(t^π_ji)
20:     end for
21:     π(i) = arg max_b u^π_ib
22: end for
```

4.2.3 Algorithm

Since both components of our CABS problem are hard, we must carefully choose the interaction between these components in our solution. Unlike the beam selection problem, the client association problem, although hard, can be solved more efficiently. Hence, ProBeam proposes and employs a simple but efficient client association algorithm as the core building block for solving the CABS problem. At a high level, it solves the client association problem for a given beam configuration and the resulting utility is used to manipulate the beam configuration of small cells in an iterative manner till an efficient CABS solution is attained. The algorithm is given in Algorithm CABS.

The input to the algorithm is the average client SNR (ρ_{ji}^b) for the epoch with respect to its neighboring small cells when they employ different beams ($b \in \mathcal{B}$) in isolation (step 1). Using the approach in Section 4.1, the CC can then determine the average client rates in the presence (r_{ji}^{π}) and absence (r_{ji}^b) of interference. The CC first determines a bootstrap beam configuration for the small cells as follows (steps 3-12). For each of the small cells, it determines the beam that yields the highest utility in the absence of interference, assuming all active clients can be potentially associated with it, i.e., $\pi(i) = \arg\max_{b \in \mathcal{B}} \{\sum_{j \in \mathcal{K}} x_{ji} U(t_{ji}^b)\}$. Note that t_{ji}^b depends on the scheduling policy and is hence coupled with the set of clients associated with the small cell. For example, in proportional fairness, $t_{ji}^b = \frac{NMw_j r_{ji}^b}{\sum_{k \in \mathcal{K}} x_{ki} w_k}$. Hence, even to determine a beam initialization $\pi(i)$, one needs to determine the set of clients (x_{ji}) that maximize the utility for the given beam in the absence of interference.[1] This can be done optimally (easy to verify) by adding users one by one such that incremental utility is maximized (steps 6-10). Specifically, for proportional fairness, the incremental utility (step 7) would correspond to,

$$j^* = \arg\max_{j \in \mathcal{K} \setminus \mathcal{L}} \sum_{k \in \mathcal{L} \cup j} w_k \log(\frac{NMr^b_{ki}}{1+|\mathcal{L}|}) - \sum_{k \in \mathcal{L}} w_k \log(\frac{NMr^b_{ki}}{|\mathcal{L}|})$$

After the beam initialization, CABS algorithm perturbs the beam choice for each of the small cells, one by one and one beam at a time. For each of the beam choices at a given cell ($\pi(i) = b$), CABS retains the rest of the beam choices for the other cells unchanged and solves the client association problem for all the small cells jointly under the updated beam configuration to determine the new utility (steps 16-19). CABS then fixes the beam choice for the small cell as the one that yields the highest utility among all its choices (step 21). The same process is repeated for updating the beam choice for each of the small cells sequentially (steps 14-22). Note that, although after one complete round of beam updates for each of the small cells (along with joint client re-association), we cannot guarantee convergence to the optimal solution, our evaluations in Section 5 reveal this is sufficient to obtain a performance very close to that of exhaustive search for beam configurations. CABS runs in $O(|\mathcal{K}|^2|\mathcal{S}|^2|\mathcal{B}|)$, with a large portion of the complexity coming from the client association module $O(|\mathcal{K}|^2|\mathcal{S}|)$.

4.2.4 *Performance Guarantee*

Given the hardness of the joint CABS problem, it is hard to establish an approximation guarantee for the entire algorithm. However, we can establish the following performance guarantee for the core building block in CABS, namely the client association part when the popular proportional fair scheduling policy is considered at the small cells.

THEOREM 2. *CABS is a $\frac{1}{2}-$approximation algorithm under proportional fairness when beam configuration is given.*

We provide some definitions on matroid and sub-modularity that are relevant for the proof.

Partition Matroid: Consider a ground set Ψ and let S be a set of subsets of Ψ. S is a matroid if, (i) $\emptyset \in S$, (ii) If $P \in S$ and $Q \subseteq P$, then $Q \in S$, and (iii) If $P, Q \in S$ and $|P| > |Q|$, there exists an element $x \in P \backslash Q$, such that $Q \cup \{x\} \in S$. A partition matroid is a special case of a matroid, wherein there exists a partition of Ψ into components, $\phi_1, \phi_2, \ldots$ such that $P \in S$ if and only if $|P \cap \phi_i| \le 1,\ \forall i$.

Sub-modular function: A function $f(\cdot)$ on S is said to be sub-modular and non-decreasing if $\forall x, P, Q$ such that $P \cup \{x\} \in S$ and $Q \subseteq P$ then,

$$\begin{aligned} f(P \cup \{x\}) - f(P) &\le f(Q \cup \{x\}) - f(Q) \\ f(P \cup \{x\}) - f(P) &\ge 0, \quad \text{and } f(\emptyset) = 0 \end{aligned}$$

PROOF. The sub-optimality of maximizing a sub-modular function over a partition matroid using a greedy algorithm of the form $x = \arg\max_{x \in \phi_i} f(P \cup \{x\}) - f(P)$ in every iteration was shown to be bounded by $\frac{1}{2}$ in [14]. We will now show that CABS is such an algorithm (step 18 being the key step), with our client association objective for a given beam configuration (π) corresponding to a sub-modular function to obtain the desired result.

[1]Note that accommodating all users can hurt the utility due to fixed frame resources but varying client rates.

Consider the ground set to be composed of the following tuples.

$$\Psi = \{(i,j) : \ i \in [1 : |\mathcal{S}|] \cup \emptyset, j \in [1 : |\mathcal{K}|]\}$$

Now Ψ can be partitioned into $\phi_j = \{(i,j) : i \in [1 : |\mathcal{S}|] \cup \emptyset\}, \forall j$. $i = \emptyset$ allows for the possibility of clients not being scheduled in an epoch. Let R be defined on Ψ as a set of subsets of Ψ such that for all subsets $P \in R$, we have (i) if $Q \subseteq P$, then $Q \in R$; (ii) if element $x \in P \backslash Q$, then $Q \cup \{x\} \in R$; and (iii) $|P \cap \phi_j| \le 1,\ \forall j$. This means that R is a partition matroid. Now, it is easy to see that any $P \in R$ will provide a feasible schedule with at most one feasible association to a small cell for each client ($|P \cap \phi_j| \le 1,\ \forall j$), thereby allowing the partition matroid R to capture our client association problem. Since each client can associate to only one small cell, our client association objective can be given as,

$$\begin{aligned} f(P) &= \sum_{i \in \mathcal{K}} \mu_i(P) \\ \text{where, } \mu_i(P) &= \sum_{j:(i,j) \in P} w_j \log(\frac{NMw_j r^\pi_{ij}}{\sum_{k:(i,k) \in P} w_k}) \end{aligned}$$

It can be seen that if $Q \subseteq P$, then $\mu_i(Q) \le \mu_i(P)$ since the algorithm picks only elements that result in positive incremental utility. Hence, it only remains to be shown that for an element (i, ℓ) such that $P \cup \{(i, \ell)\}$ forms a valid schedule, then $f(P \cup \{(i, \ell)\}) - f(P) \le f(Q \cup \{(i, \ell)\}) - f(Q)$. Now, define incremental utility $\Delta_P(i, \ell) = f(P \cup \{(i, \ell)\}) - f(P)$ and similarly define $\Delta_Q(i, \ell)$. Applying the objective function and simplifying, we can show that,

$$\begin{aligned} \Delta_P(i,\ell) &= w_\ell \log(NMw_\ell r^\pi_{i\ell}) - w_\ell \log(w_\ell + \sum_{k:(i,k) \in P} w_k) \\ &\quad - \sum_{j:(i,j) \in P} w_j \log(\frac{w_\ell + \sum_{k:(i,k) \in P} w_k}{\sum_{k:(i,k) \in P} w_k}) \\ \Delta_Q(i,\ell) &= w_\ell \log(NMw_\ell r^\pi_{i\ell}) - w_\ell \log(w_\ell + \sum_{k:(i,k) \in Q} w_k) \\ &\quad - \sum_{j:(i,j) \in Q} w_j \log(\frac{w_\ell + \sum_{k:(i,k) \in Q} w_k}{\sum_{k:(i,k) \in Q} w_k}) \end{aligned}$$

Thus, the difference between $\Delta_P(i, \ell)$ and $\Delta_Q(i, \ell)$ arises in the second (reduction) term, which increases with the number of elements in the allocation thus far. Since $Q \subseteq P$, the reduction term is more for P than for Q, resulting in $\Delta_P(i, \ell) \le \Delta_Q(i, \ell)$. This establishes that the function $f(P)$ is indeed sub-modular. Further, our client association problem aims to maximize this non-decreasing sub-modular function over a partition matroid. Hence, picking the (client, small cell) pair yielding the highest marginal utility for a given beam configuration in CABS (steps 16-19) would correspond to determining

$$(i^*, j^*) = \arg\max_{(i,j) \in R} \{f(P \cup \{(i,j)\}) - f(P)\}$$

Thus, the sub-optimality of $\frac{1}{2}$ would then follow from the result in [15]. □

4.3 Scheduling

Once the CC determines the beam configuration and client association for the epoch, the appropriate beam and allowable client set are notified to each of the small cell BSs for configuration. Each small cell BS then locally runs its scheduling algorithm (e.g., proportional fair) among the associated clients for each frame in the epoch, while employing the chosen beam for its transmissions. Further, instantaneous channel rate feedback from clients is used in per-frame scheduling for leveraging multi-user diversity.

4.4 Practical Considerations

Mobile clients: While beamforming algorithms work well for static clients, it is important to understand their limitations with respect to mobile clients. Note that, any adaptive beamforming scheme that relies on fine grained channel state information (CSI) will be highly sensitive to lack of timely and accurate CSI, both of which are hard to obtain during mobility. On the other hand, switched beamforming relies only on coarse grained channel feedback (SNR or RSSI) and hence is less sensitive to mobility. As long as the epoch duration is not long enough (several seconds is reasonable), pedestrian to moderate vehicular speeds can be accommodated without warranting a completely new beam to be employed for the client.

Epoch duration: Keeping the epoch duration long is conducive for implementation and overhead. However, it must also be capable of tracking traffic dynamics and client mobility. Allowing for a few seconds of epoch duration strikes a good balance between these objectives.

5. SYSTEM EVALUATION

Testbed and prototype implementation: Our WiMAX testbed consists of four small cells (deployed in an indoor enterprise environment), clients and a central controller as depicted in Figure 7. The small cell BS is a PicoChip [16] WiMAX platform based on IEEE 802.16e standard [17]. The BS is tuned to operate in a 10 MHz bandwidth with the center carrier frequency of 2.59 GHz, for which we have obtained an experimental license to transmit WiMAX signals over the air. In the absence of a macro cell to coordinate with, we use a GPS module to synchronize the WiMAX frame transmissions across the small cells. Each BS has an eight element (analog) phased array antenna [18] connected to its RF port. The antenna array generates sixteen overlapping beam patterns of 45° each, spaced 22.5° apart to cover the entire azimuth of 360°. The BS controls the antenna array through a serial port application that we have developed in C. There is a delay of one frame (5 msec) before a particular beam pattern is actually applied by the antenna following the command from the application. This is not an issue given the time scale of epoch or the measurement phase.

ProBeam is standards compatible and works with commercial off-the-shelf clients. We use Windows laptops with a WiMAX interface [19] and omni-directional antennas as our clients. Investigating directionality at the clients is part of our future work. We select 30 locations as marked in Figure 7 for client deployments. The clients are oblivious to beam selection at BS and simply measure the SNR and report them back to the BS for SINR estimations through standard feedback mechanisms. Our experiments have verified that the SNR received on each beam is relatively stable over several seconds for static clients. This gives confidence to the SNR measurements reported by clients in the measurement phase.

All algorithms (CABS and reference schemes) are implemented on the CC and do not require any changes or operational overhead to the BS. All BSs are connected to the CC through an ethernet switch in our set-up.

5.1 Prototype Evaluations

Topologies and rate adaptation: Each data point in our result is averaged over multiple topologies, which are generated by picking random subsets of client locations (among 30) for a given number of clients. Further, unless otherwise specified, we consider topologies with four small cells and twenty clients. To remove the influence of rate adaptation algorithms, we consider an ideal PHY rate adaptation by trying out all MCS and record the highest throughput (best MCS) for a client given a network configuration.

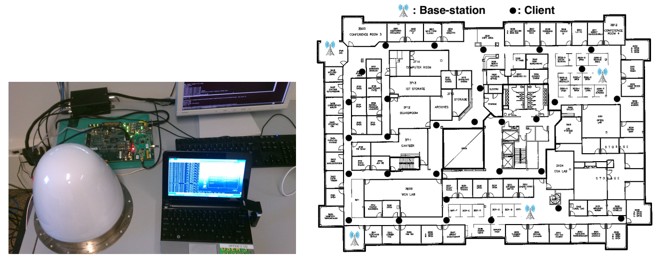

Figure 7: Small cell, beamformer, client and deployment

Reference schemes: We evaluate the performance of our CABS algorithm in ProBeam against the following benchmark algorithms.

- *Decoupled:* Client association is decoupled and first computed based on SNR, followed by determination of coordinated beams for each BS using the same beam selection component as in CABS.
- *CABS-all:* Allows for joint determination of client association and beam selection as in CABS but requires that all clients be associated and scheduled in every epoch.
- *UB-beam:* Employs the same client association component as in CABS but exhaustively searches over all possible beam combinations at BSs - serves as an upper bound for beam selection in CABS.
- *UB-assoc:* Employs the same beam selection component as in CABS but exhaustively searches over all possible combinations of client association.

Evaluation metrics: We consider the following metrics.

- *Throughput:* Aggregate throughput of all clients in the network.
- *Utility:* Captures both throughput and fairness; aggregate utility of all clients: $\sum_{j \in K} w_j \log(T_j)$ (details in subsection 4.2).
- *Fraction of scheduled clients:* Captures the number of clients not scheduled in an epoch to improve spatial reuse (in CABS and upper bounds).
- *Load balancing factor:* Measures Jain's fairness index among the number of clients associated with each BS.

Throughput: Figure 8(a) presents the throughput results as a function of number of clients in the network. Three observations can be made: (i) CABS' performance is within 96% of that of exhaustive beam search and is not impacted by client density. Given the complexity of the latter, CABS provides a fine balance between performance and complexity. (ii) The increased spatial reuse from jointly addressing client association with beamforming (CABS-all) provides gains as high as 50% (over the decoupled approach). Further, the gains are more pronounced at higher client density, where it becomes harder to isolate interference between small cells without a joint optimization that allows for flexible client association. (iii) Interestingly, by going one step further and allowing some clients from not being scheduled in a given epoch provides CABS with an additional 50% gain over CABS-all, resulting in *a net gain*

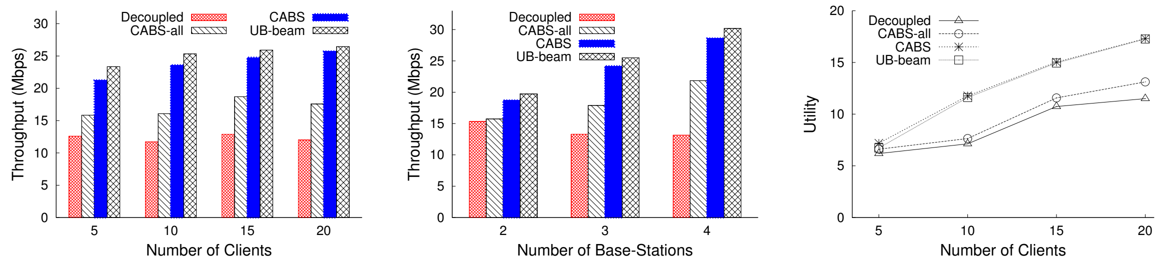

(a) Throughput (b) Throughput (c) System utility

Figure 8: Experimental evaluation of ProBeam with 4 small cells.

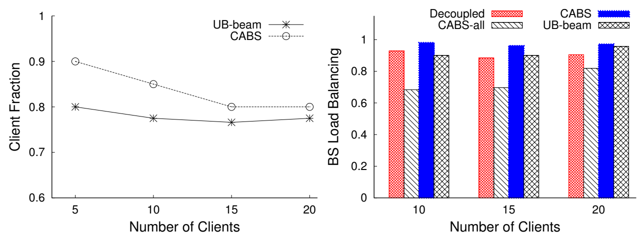

(a) Scheduled client fraction (b) Network load balancing

Figure 9: Effective client management in ProBeam.

of around 115% over the decoupled approach. Removing even a small fraction of bottleneck clients from scheduling in an epoch can greatly improve the spatial reuse configuration between small cells.

The impact of interference from increased number of BSs is presented in Figure 8(b). The ability to jointly address client association with beam selection helps CABS handle interference effectively, the benefits of which are more pronounced with larger number of interferers.

Fairness: Recall that some of the reuse gains in CABS comes from removing a subset of clients from scheduling in a given epoch. While starvation of such clients is avoided across epochs, it is important to understand if the throughput gains of CABS are not realized at the expense of fairness even within an epoch. The utility measure helps account for fairness within an epoch, whose results are presented in Figure 8(c). It can be clearly seen that CABS' utility is very close to that of its upper bound and outperforms that of the (baseline) decoupled approach. Thus, *adopting a utility based approach to joint CABS, enables ProBeam to bypass some clients from an epoch to maximize reuse gains without compromising on fairness*.

Note that if the number of clients bypassed is large, this would automatically reflect in a reduced system utility. Hence, to further verify this, we present the fraction of scheduled clients in an epoch in Figure 9(a). This clearly shows that only a small fraction of clients (10-20%) are bypassed in CABS. The upper bound is more aggressive in deferring clients to the next scheduling epoch, which in turn contributes to its marginal throughput gains over CABS (Figure 8(a)).

Load balancing: A by-product of *utility maximization in CABS is that it should automatically lead to load balancing*. This is because, given a fixed amount of frame resources, balancing number of users across cells, provides more resources per user and hence better aggregate utility. The load balancing factor, captured through Jain's fairness index between number of clients associated with small cells, is presented in Figure 9(b). CABS provides very good load balancing as expected. The decoupled approach does not implicitly account for load balancing, but a uniform distribution of clients automatically provides reasonable load balancing, when SNR-based client association is employed. The interesting observation is that CABS-all's load balancing suffers, especially when the number of clients is not high. Recall that CABS-all's throughput gain (over the decoupled approach) from better interference suppression (and hence reuse) through flexible association, comes at the expense of potential load imbalance across cells, especially when all clients are accommodated.

5.2 Trace-driven Simulations

Our experimental set-up with few tens of clients and three dominant interferers constitutes a realistic set-up for a cluster of small cells. However, to further understand CABS's effectiveness in much denser deployments (10 BSs and 90 clients), we resort to trace based simulations. We collect SNR traces for clients from our experimental network, feed it into a simulator running ProBeam (SINR estimation and CABS) to evaluate the various algorithms. We place our four BSs in various other locations to emulate more small cell BSs and measure SNR traces at the clients from them on all beams. Similarly, we also vary the client locations to emulate a larger set of clients and obtain corresponding SNR traces. Given the traces, we can generate a topology with a specific number of BSs and clients, by sampling BSs and clients randomly from our SNR trace database.

Our simulation results are presented in Figure 10, where throughput is measured as a fraction of that achieved by the upper bound (UB-beam). The trends in these large scale results, including the magnitude of gains possible with CABS, are very similar to those from the experiments, thereby reinforcing our inferences from the prototype evaluation. Hence, in the interest of space, we do not discuss them further. CABS close performance with respect to its upper bound in these results indicates the efficiency of its beam selection component as both the schemes employ the same client association mechanism. Given the hardness of computing a tight upper bound for the joint CABS solution, we now evaluate the efficiency of its client association component as well. We compare it against an upper bound for client association (UB-assoc) that exhaustively searches over all possible client associations, while employing the same beam selection mechanism as in CABS. The results in Figure 11 indicate that, while the sub-optimality of CABS' client association component can at most be within half of the optimal (see Sec.4.2.4) in the worst case, in practice, it yields a performance that is very close to its upper bound. Thus, *the high effi-*

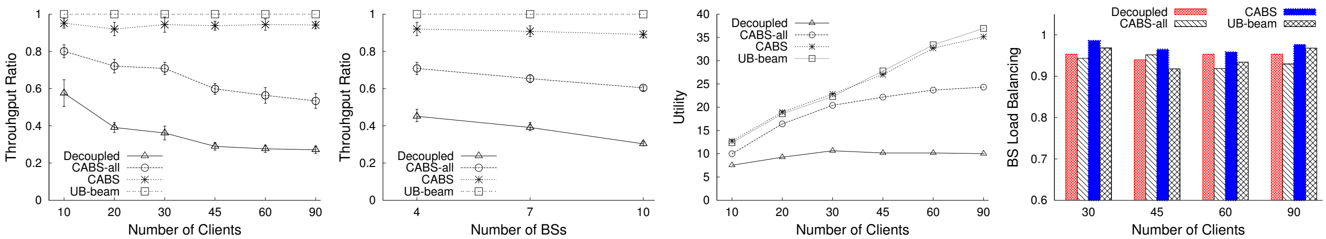

(a) Throughput ratio (b) Throughput ratio (c) System utility (d) BS load balancing

Figure 10: Large scale evaluation of ProBeam through trace-driven simulations.

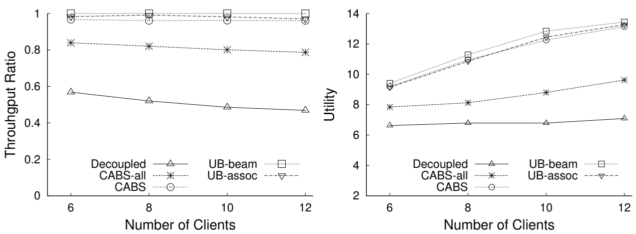

(a) Throughput ratio (b) Utility

Figure 11: Evaluation of client association component in ProBeam.

ciency of the individual components in CABS in turn synergistically contribute to the net gains seen by it.

6. CONCLUSIONS

We design and implement *ProBeam* – a practical system for improving spatial reuse through beamforming in OFDMA based small cell networks. We show that decoupling beamforming from client scheduling is necessary for practical feasibility. Further, we highlight the need to jointly address client association with beamforming to maximize the reuse benefits from the latter. ProBeam incorporates a low complexity, highly accurate SINR estimation module with less than 1 dB error ($\leq 5\%$) to determine interference dependencies between small cells. It also houses an efficient, low complexity joint client association and beam selection algorithm for the small cells that yields close-to-optimal performance. Prototype implementation in a real WiMAX networks of four small cells shows 115% of capacity gain compared to other baseline reuse schemes. We also demonstrate the scalability and efficacy of our system in larger scale settings through simulations. Most of our system components are also applicable to LTE and LTE-A with minor modifications. As part of future work, we plan to investigate synthesis of new beam patterns for beamforming based on client feedback in lieu of a pre-determined set (code-book).

Acknowledgments

We would like to thank the anonymous reviewers for their invaluable feedback. Jongwon Yoon and Suman Banerjee have been supported in part by the following grants of the US National Science Foundation: CNS-1040648, CNS-0916955, CNS-0855201, CNS-0747177, CNS-1064944, and CNS-1059306.

7. REFERENCES

[1] R. Van Nee and R. Prasad, "OFDM for Wireless Multimedia Communications", *Artech House*, 2000.

[2] M. Y. Arslan, J. Yoon, K. Sundaresan, S. V. Krishnamurthy, and S. Banerjee, "FERMI: A FEmtocell Resource Management System for Interference Mitigation in OFDMA Networks", In *ACM MobiCom*, 2011.

[3] J. Yoon, M. Y. Arslan, K. Sundaresan, S. V. Krishnamurthy, and S. Banerjee, "A Distributed Resource Management Framework for Interference Mitigation in OFDMA Femtocell Networks", In *ACM MobiHoc*, 2012.

[4] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste, "DIRC: Increasing Indoor Wireless Capacity Using Directional Antennas", In *ACM Sigcomm*, 2009.

[5] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste, "Pushing the Envelope of Indoor Wireless Spatial Reuse using Directional Access Points and Clients", In *ACM MobiCom*, 2010.

[6] Femtocells Core Specification, WMF-T33-118-R016v01.

[7] M. Y. Arslan, J. Yoon, K. Sundaresan, S. Krishnamuthy, and S. Banerjee, "Characterization of Interference in OFDMA Femtocell Networks", In *IEEE Infocom*, 2012.

[8] J. Yun and K. G. Shin, "CTRL: A Self-Organizing Femtocell Management Architecture for Co-Channel Deployment", In *ACM MobiCom*, 2010.

[9] K. Sundaresan and S. Rangarajan, "Efficient Resource Management in OFDMA Femto Cells", In *ACM MobiHoc*, 2009.

[10] S. He, Y. Huang, L. Yang, A. Nallanathan, and P. Liu, "A Multi-Cell Beamforming Design by Uplink-Downlink Max-Min SINR Duality", *IEEE Transactions on Wireless Communications*, vol. 11, no. 8, pp 2858-2867, 2012.

[11] H. Dahrouj, and W. Yu, "Coordinated Beamforming for the Multicell Multi-Antenna Wireless System", *IEEE Transactions on Wireless Communications*, vol. 9, no. 5, pp 1748-1759, 2010.

[12] E. Aryafar, A. Khojastepour, K. Sundaresan, S. Rangarajan, and E. Knightly, "ADAM: An Adaptive Beamforming System for Multicasting in Wireless LANs", In *IEEE Infocom*, 2012.

[13] ProBeam technical report, http://www.cs.wisc.edu/∼yoonj/ProBeam-TR.pdf/

[14] L. Fleischer, M. X. Goemans, V. S. Mirrokni, and M. Sviridenko, "Tight Approximation Algorithms for Maximum General Assignment Problems", In *ACM-SIAM symposium on Discrete algorithm*, 2006.

[15] M. Fisher, G. Nemhauser, and G. Wolsey, "An Analysis of Approximations for Maximizing Submodular set Functions-II", *Mathematical Programming Study*, 1978.

[16] PicoChip Femtocell Solutions, http://www.picochip.com/

[17] IEEE 802.16e-2005 Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems, *IEEE 802.16e standard*.

[18] Fidelity Comtech, http://www.fidelity-comtech.com/

[19] Accton Wireless Broadband, http://www.awbnetworks.com/