

Experimental Characterization of Interference in OFDMA Femtocell Networks

Mustafa Y. Arslan* Jongwon Yoon† Karthikeyan Sundaresan‡ Srikanth V. Krishnamurthy* Suman Banerjee†

*UC Riverside: {marslan, krish}@cs.ucr.edu †U Wisconsin Madison: {yoonj, suman}@cs.wisc.edu ‡NEC Labs America Inc: karthiks@nec-labs.com

Abstract—The increase in mobile data usage is pushing broadband operators towards deploying smaller cells (*femtocells*) and sophisticated access technologies such as OFDMA. The expected high density of deployment and uncoordinated operations of femtocells however, make interference management both critical and extremely challenging. Femtocells have to use the same access technology as traditional macrocells. Given this, understanding the impact of the system design choices (originally tailored to well-planned macrocells) on interference management, forms an essential first step towards designing efficient solutions for next-generation femtocells. This in turn is the focus of our work. With extensive measurements from our WiMAX OFDMA femtocell testbed, we characterize the impact of various system design choices on interference. Based on the insights from our measurements, we discuss several implications on how to efficiently operate a femtocell network.

I. INTRODUCTION

To satisfy the ever-increasing mobile data needs of users, broadband operators have recently started to deploy femtocells with OFDMA [1]. They are installed in homes and enterprises, and use the same spectrum and access technology as macrocells, while connecting to the core network through cable or DSL backhaul. Femtocells, being deployed indoors, present numerous advantages to mobile broadband customers and service providers. For example, the users can continuously experience the “5-bar” cellular coverage in areas where the macrocell signal is poor. In addition, the small range and low cost of femtocells offer both cost-effective ways of providing coverage and increased system capacity via spatial reuse.

While the above advantages motivate femtocells, they tend to get deployed in an unplanned manner by end users, without direct coordination of the cellular operators. Hence, they inevitably lead to equivalent unplanned interference. If not properly understood, evaluated and accounted for, this can cause significant degradation in femtocell performance.

Femtocells have to inter-operate with and use the same access technology as macrocells. Although femtocells inherit the access method from macrocells, they cannot adopt the same interference mitigation solutions for two reasons. (i) Femtocell deployments are dense and uncoordinated unlike planned deployments of macrocells. Thus, unlike in macrocells where interference is localized at cell-edges, interference is less predictable and more pervasive across femtocells. (ii) Unlike in macrocells, it is harder to achieve synchronization across interfering femtocells deployed by different users. The impact of this lack of synchronization on interference management is not well understood; one might however, expect that this would make interference management harder.

In this paper, we present one of the first measurement studies of OFDMA femtocells, conducted using our WiMAX testbed. Since femtocells have to inter-operate with macrocells, solutions to cope with interference among femtocells cannot be

designed from scratch and have to strictly follow the cellular access standards. The effect of the design choices made in the macrocellular context (i.e., the system parameter options defined in the standards) have to be re-evaluated in femtocell deployments. We provide an in-depth understanding of the impact of these design choices, on interference management in femtocells. The design choices are categorized along three broad dimensions. (i) Resource isolation across femtocells: isolation of transmissions between interfering cells can be either for the data payload alone or for both the data and control payload. (ii) Resource mapping (sub-channelization) in each femtocell: OFDMA sub-channels that provide the basic granularity of transmissions at the MAC layer, consist of a group of sub-carriers (tones) at the PHY layer. Varying the sub-carrier composition (grouping) of sub-channels using (a) contiguous or (b) distributed grouping, can impact performance. (iii) Synchronization: lack of tight synchronization between femtocells can exacerbate interference given the rigid frame structure (designed for macrocells).

Our study reveals several interesting effects. We carefully examine these and provide inferences that help understand the implications of the aforementioned design choices in the interference-rich femtocell environment. Our key findings are:

- Resource isolation for the data payload alone helps alleviate interference in most of the cases. Further, its degree of isolation is not limited as opposed to isolation for both the control and the data payload.
- A contiguous composition of sub-carriers at the PHY layer, reduces interference that arises from frequency offsets across cells.
- Frame-level synchronization between base-stations (BSs) yields predictable benefits from resource isolation. More interestingly, there is still merit to isolating resources even in the absence of synchronization.

The rest of the paper is organized as follows. In §II, we overview WiMAX and OFDMA, and discuss related work. §III describes our experimental platform and methodology. §IV, V and VI describe our measurements on resource isolation, resource mapping and synchronization. The inferences from our measurements and guidelines for efficient femtocell network operations are in §VII. Our conclusions form §VIII.

II. RELATED WORK AND BACKGROUND

Related Work: Although solutions (e.g., [2], [3]) exist to address interference in OFDMA macrocells, they cannot be reused in the femtocell context. Since femtocells lack the desired features of localized interference, planned deployments and coordinated operations of macrocells, they require novel resource management solutions for interference mitigation. There are recent studies that propose such solutions [4], [5],




Fig. 1. Illustration of sub-channel composition in WiMAX. Two (out of 30) sub-channels are shown for clarity.




Fig. 2. The testbed deployment (a), resource isolation choices that we consider for femtocells (b) and FFR in macrocells (c).

[6]. However, [4] and [5] are theoretical studies with several simplifying assumptions that prevent their adoption in practice. While we designed and implemented a practical resource management solution in [6], we did not explore the design space and several assumptions were made with respect to system parameters (e.g., femtocells were manually synchronized). In this regard, several questions such as “*Is the solution in [6] still viable with lack of synchronization?*” still remain to be answered. Here, we undertake a comprehensive study to answer several other questions and understand the design choices that influence interference management in femtocell deployments.

WiMAX Preliminaries: A WiMAX frame (see [6] for an illustration) is a two-dimensional template that carries data to multiple mobile stations (MSs / clients) across both time (symbols) and frequency (sub-channels). The combination of a symbol and a sub-channel constitutes a *tile*, which is the basic unit of resource allocation at the MAC layer. Data to MSs are allocated as rectangular bursts of tiles in a frame. In addition to data, the frame includes the preamble (used by the MSs to associate to a BS) and the control payload. To decode their data, the MSs first decode the control payload, which contains parameters such as the modulation and coding scheme (MCS) used.

While the frame allocates data at the granularity of tiles, the bits map to sub-carriers at the PHY layer. WiMAX offers three options for the grouping of sub-carriers to form a sub-channel (see Fig. 1). **FUSC (full usage of sub-carriers):** sub-carriers composing a sub-channel are picked in a completely distributed manner; **PUSC (partial usage of sub-carriers):** contiguous sub-carriers are first grouped into clusters and distributed clusters are then grouped to form a sub-channel; **AMC (adaptive modulation and coding):** contiguous set of sub-carriers are grouped to form a sub-channel. PUSC is the mandatory default mode in 802.16e. The first generation of WiMAX devices implement PUSC and FUSC alone. Hence, we restrict our measurements to these modes.

Our Work in Perspective: To understand how to cope with interference in femtocells, we conduct extensive measurements on a WiMAX femtocell testbed. We explore several design choices and seek answers to the following:

- Is it sufficient to isolate resources (i.e., tiles) for data bursts alone or is isolation also needed for the control part of a frame?

- What is the impact of different sub-carrier grouping modes (resource mapping) on resource isolation?
- What is the impact of the lack of synchronization across interfering BSs, on resource isolation and mapping?

III. SYSTEM DESCRIPTION

Experimental Setup: Our testbed consists of three femtocells deployed in an indoor enterprise environment (Fig. 2(a)), each with one client. We use PicoChip’s femtocells that run 802.16e (WiMAX). Our clients are commercial USB WiMAX dongles (attached to laptops) from Accton. The cells have 8.75 MHz of bandwidth with the carrier frequency of 2.59 GHz. For this frequency, we have an experimental license from FCC to transmit WiMAX signals.

Synchronous Transmissions: OFDMA uses synchronous channel access where a BS transmits a downlink frame *periodically* (5 ms for WiMAX). Although frames are synchronized between a client and its BS, they may not be synchronized across BSs. We perform experiments with both synchronized and unsynchronized frames among the BSs. While each BS has its own internal clock, they are initially not synchronized with each other. To synchronize the cells, we use external GPS modules (antennas placed near a window) from TeraSync [7] to provide a 1 pps (pulse per second) signal to each BS, thereby aligning the start times of downlink frames.

Experimental Methodology: We focus on cell 1 (with client 1) for our measurements. The other two BSs project interference (unless noted otherwise) on cell 1. Here, even switching on other BSs (without any data transfer) will cause interference, since a BS transmits the preamble and control payload even when no clients are associated with it. However, with data transfer, the impact of interference is more severe. To understand these different aspects of interference, we pick client 1 and measure its throughput with either one or two other interfering cells, both with and without data transfers in the interfering cells (to clients 2 and 3). We consider saturated downlink UDP traffic from the BSs to the clients generated by *iperf*. In this setting, interference is due to transmission in cells 2 and 3, on a set of tiles that collide with that of the BS - client pair under consideration.

Link Adaptation: The impact of interference on a client depends on the interference power received (i.e., location) and the MCS used for the transmission. The SNR required for correct decoding increases with higher MCS levels and




Fig. 3. The performance of different resource isolation strategies with interference.

thus, the impact of a given interference power is more severe with a higher MCS. Note that femtocells are meant to be user-installed at homes and in enterprises. Hence, the baseline strategy is one where a femtocell operates on its entire spectrum, while performing link adaptation (MCS selection) for each of its clients. To eliminate the artifacts of any particular link adaptation algorithm, we adopt an ideal link adaptation strategy, whereby for each data point, we sequentially run the experiment over all MCS levels and record the one delivering the highest throughput for the given topology. Since we experiment in a slowly-varying static indoor environment, it is feasible to consider link adaptation as above.

We call a given triplet $\{bs, cl, int\}$ a *topology* where bs is the BS that the client is associated with, cl is the location of the client and int is the set of BSs that project interference on the client. Each measurement point corresponds to a topology and is obtained by running an experiment for 7 minutes (one minute for each MCS), measuring the throughput and averaging it over several such runs.

We categorize our study along the following three dimensions: (i) resource isolation, (ii) resource mapping, and (iii) synchronization. To isolate the benefits of resource isolation and mapping from the impact of synchronization, we consider the first two dimensions with synchronized femtocells.

IV. RESOURCE ISOLATION

In this section, we examine the impact of MAC level resource isolation choices on interference. Resource isolation refers to assigning orthogonal sets of tiles to cells that interfere with each other. As shown in [6], this is achieved by assigning orthogonal sets of sub-channels to each cell (i.e., frequency domain isolation or FDI). There are three strategies with respect to isolation: (a) the baseline strategy where each BS uses all the sub-channels in the spectrum (i.e., no isolation), (b) FDIdata, which assigns orthogonal sub-channels only for the data payload (see Fig. 2(b)) and (c) FDICtrl, which uses orthogonal sub-channels for both the control and data payload.

Cellular standards typically restrict the degree of isolation for FDICtrl. For example, macrocells use 1:3 fractional frequency reuse or FFR (see Fig. 2(c)), where each cell is partitioned in to three sectors (labelled A, B, C). To support 1:3 FFR, the frame is divided to form three orthogonal segments of sub-channels (each sector uses one segment). If there are more than three cells in a contention region (which is likely in a dense femtocell deployment), some segment is inevitably shared by multiple cells for the control payload (while the data can be partitioned). We seek to understand “*What is the impact of not isolating the control part and whether data isolation can still provide benefits in such scenarios?*”

We consider two cells (cells 1 and 2) and measure the throughput at client 1 with interference, with three strategies: baseline, FDIdata and FDICtrl. The histogram of the results

over several topologies is presented in Fig. 3(a). The median throughputs of the baseline, FDIdata and FDICtrl are around 0.3, 2.5 and 5 Mbps, respectively. This indicates that while complete isolation of both control and data parts can significantly alleviate interference, isolating sub-channels for the data part alone (with control parts experiencing interference) can still deliver significant benefits over the baseline strategy.

Next, we investigate the robustness of FDIdata and FDICtrl to increased interference. We now include cell 3 with the same set of sub-channels as cell 2 (already orthogonal to those in cell 1); thus, it projects additional interference on the control payload for client 1. We present a subset of our results in Fig. 3(b). We see that increased interference does not have an impact on FDICtrl due to complete isolation (control + payload). However, increased interference on the control payload does degrade performance with FDIdata by 20-30%.

Our next experiment is with two cells (cell 1 and 2) and their clients. Initially both clients are close to their BSs (see Fig. 3(c)). Gradually, we increase the distances of both clients from their BSs and move them closer to the interfering BS (as indicated by the arrows). The corresponding aggregate throughput for the three strategies shown in Fig. 3(d), clearly reinforce the benefits with FDIdata. We see that if the clients are close to the BS, the difference in throughputs with FDICtrl and FDIdata is not significant since the control part is not susceptible to interference; however, as they move away, the control part is affected and the throughput with FDIdata decreases (still outperforms baseline). In summary, *although perfect isolation of both data and control parts is best, if control part isolation is not feasible, data part isolation can still deliver significant throughput gains.*

A Microscopic View: Next, we provide some microscopic insights from our experiments. Our clients’ connection manager reports statistics including signal strength, erroneously decoded packets and the cause of errors. The causes of errors prove especially helpful in analyzing our results. Although both the control and data payloads are isolated in FDICtrl, preamble collisions inevitably occur. We observed rare cases with FDICtrl where, the MS disconnects from its BS due to interference. We find that placing the MS at a location very close to the interfering BS, causes it to drop the connection since it cannot decode the preamble from its BS (due to high power from the interfering BS). This was also common to baseline and FDIdata. For FDIdata, overlap of the control payloads across cells is also a factor. We found that at locations where the control payload was corrupted, the client did not receive any packets (0 Mbps as in the 25m point on Fig. 3(d)). This is because of the loss of DL-MAP which contains the parameters for correct decoding of the data bursts. Although the data bursts are isolated from interference, the corruption of DL-MAP impacts performance.




Fig. 4. Impact of Sub-carrier Contiguity on Interference.

V. RESOURCE MAPPING

We have so far considered isolating resources at the MAC (logical) level of sub-channels to alleviate interference. However, the sub-carrier composition of sub-channels remained the same across cells. Next we investigate how varying the sub-carrier composition at the PHY, can impact performance.

Contiguous vs. Distributed Grouping: Recall from §II that current WiMAX implementations widely adopt two modes of sub-carrier grouping: contiguous (PUSC) and distributed (FUSC). Before considering interference, we experiment with PUSC and FUSC without interference to understand their baseline performance. We present client 1’s throughput at various locations in Fig. 4(a). Note that FUSC uses 768 sub-carriers, while PUSC uses only 720 per symbol. Hence, a normalizing line is drawn at $y = \frac{768}{720}x = \frac{16}{15}x$. The clustering of points around this line indicates that the throughput of PUSC and FUSC are similar in the absence of interference.

We now introduce interference from cell 2 and measure client 1’s throughput at various locations. While PUSC can isolate both the data and control parts, FUSC does not allow for control part isolation [8]. With PUSC, we consider both FDIctrl and FDIdata. The results from a representative set of topologies are in Fig. 4(b). Three inferences are in order: (i) PUSC outperforms FUSC with an average gain of 35%. This can be attributed to the following effect. While the frequency offsets between cell 1 and its client 1 are taken care of using signaling in the frame, the synchronization provided externally between the cells does not address the frequency offsets between the interfering cell (cell 2) and client 1. This disturbs the orthogonality of sub-carriers between cell 1 and cell 2. Hence, although orthogonal sub-channels are used at the MAC, frequency offsets across cells allow energy from sub-carriers of cell 2 to spill over the adjacent sub-carriers of cell 1. This degrades the decoding quality at client 1. In PUSC, a contiguous set of 14 sub-carriers are grouped to form a cluster; clusters are then grouped to form a sub-channel. This contiguity of sub-carriers (within a cluster) in a sub-channel in PUSC, limits the corruption of sub-carriers (due to frequency offsets) only to those at the edges of the clusters. However, since the sub-carriers in a sub-channel are completely distributed in FUSC, most of the sub-carriers are vulnerable to the above effect. This explains the loss in performance. (ii) As one might expect, with PUSC, FDIctrl outperforms FDIdata. (iii) Interestingly however, in location 4, we find that FDIctrl does not perform the best. This is because, when the control part is isolated, it expands into the data part, thereby reducing the number of tiles available for the data payload. Hence, if the interference on the control part is not severe, the loss in throughput due to the expansion of the control part outweighs the benefits of FDIctrl.

A more extensive set of measurements with each strategy is in Fig. 4(c). It is seen that PUSC with FDIdata provides

a median gain of 50% over FUSC, while PUSC with FDIctrl can provide an additional gain of 50%. Further, for about 20% of the topologies where interference on the control part is such that it does not prevent correct decoding, FDIctrl does not pay off. In summary, *distributed sub-carrier grouping (FUSC) is more vulnerable to frequency offsets across cells. This degrades performance in the presence of interference. Further, FUSC can isolate resources only for the data part unlike PUSC and hence, is limited in addressing interference. In cases without interference, FUSC outperforms PUSC due to the higher number of sub-carriers.*

VI. SYNCHRONIZATION

In the experiments thus far, the femtocells were synchronized via GPS modules. There are tight synchronization requirements in macrocell OFDMA networks. In upcoming femto standards (e.g., 802.16m) provisions are being made for achieving synchronization between femtocells, potentially using self-organizing (SON) servers. However, it is unclear if *tight* synchronization can be achieved among femtocells. Hence, it is important to understand the impact of synchronization or lack thereof on resource isolation and resource mapping. Next, we consider cells 1 and 2 without frame level synchronization. With the preamble being transmitted at 3-5dB higher power, interference patterns could change, potentially causing further throughput degradation.

Impact on Resource Isolation: First, we study synchronization and resource isolation. The throughput at client 1 in the presence of interference is measured over numerous locations and its CDF is plotted in Fig. 5(a). By isolating both the control and data parts of the frame between interfering cells, FDIctrl is relatively immune to lack of synchronization and continues to outperform baseline and FDIdata. While the preamble of one cell can corrupt the data of the other cell, this corruption is restricted to one symbol with the rest of the frame being completely isolated with FDIctrl. Between baseline and FDIdata, we find that there are three distinct regions: baseline outperforms FDIdata, FDIdata outperforms baseline, and both strategies suffer.

To gain a deeper understanding, we performed a more detailed six-stage experiment with the timeline in Fig. 4(d): (1) Cell 1 alone starts operating and transferring data to client 1 on all sub-channels (All SC) from $t=0$ onwards, (2) cell 2 switches on and starts transmitting only the preamble and control from $t=30$, (3) cell 2 starts transferring data to client 2 on all sub-channels from $t=60$, (4) cell 1 and cell 2 switch to different sets of sub-channels from $t=90$ (FDIdata), (5) cell 2 stops transferring data and goes back to transmitting only the preamble and control from $t=120$, and (6) cell 2 switches off at $t=150$. Re-starting the BSs for every experimental run changes the frame alignment between cells and hence, the interference pattern. Two of the prominent patterns that repeat




Fig. 5. Impact of Lack of Synchronization on Resource Isolation and Sub-channelization.

over numerous experiments are shown in Figs. 5(b) and 5(c). In one of the patterns in Fig. 5(b), we find that even when only the preamble of cell 2 is switched on, it completely inhibits the transmission of cell 1, which does not recover until cell 2 is switched off. This corresponds to the case where both baseline and FDIdata suffer (left region in Fig. 5(a)). In the other pattern in Fig. 5(c), we find that when cell 2's preamble is switched on, it does not completely disrupt cell 1's transmission but degrades its throughput from 15 Mbps to 11 Mbps. However, when cell 2 starts transferring data on all sub-channels, this completely interferes with cell 1's transmission. However, when both the cells isolate their data in stage 4, cell 1's transmission starts recovering and approaches the maximum throughput possible with half the sub-channels. This corresponds to the case, where FDIdata outperforms baseline (middle region in Fig. 5(a)). In the remaining 40% of topologies, the interference between the cells was not severe enough to prevent decoding even with the lowest MCS. In this case, link adaptation was sufficient to address interference (both baseline and FDIdata support the same MCS) and hence, isolating resources unnecessarily reduces throughput compared to using all sub-channels. This explains the right region in Fig. 5(a), where baseline outperforms FDIdata.

Thus, we find that *even without synchronization, resource isolation between interfering cells is extremely beneficial - isolation of both data and control parts continues to deliver significant benefits, while isolating the data part alone also contributes to large gains in over 30% of the scenarios.*

Impact on Resource Mapping: Next, we study resource mapping with lack of synchronization. We measure the throughput with PUSC and FUSC for cells 1 and 2. The BSs use FDIdata and the throughput of client 1 is measured at various locations. Fig. 5(d) depicts the CDF of these measurements. PUSC outperforms FUSC in low throughput regimes, where interference dominates performance. This is consistent with our observations in §V; FUSC is still vulnerable to frequency offsets between the BSs. However, for throughput samples larger than 2.5 Mbps, FUSC outperforms PUSC. In these regimes where interference is not the dominant factor, client 1 alleviates the effects of interference with link adaptation. We also observe that such samples are clustered around 2.977 and 3.247 Mbps with PUSC and FUSC, respectively (with the best MCS being the same for both schemes). The clustering can be explained by the inherent benefit with FUSC due to the higher number of sub-carriers; the ratio corresponds to $\approx \frac{16}{15}$. In summary, we find that *PUSC is less vulnerable to frequency offsets across cells and this helps in handling interference better than FUSC, even without synchronization.*

VII. DESIGN IMPLICATIONS

Our measurements provide an understanding of how to address interference in femtocell networks. The key inferences

are summarized below.

- Resource isolation, even when achieved on only the data part, can still alleviate interference. While complete isolation (FDICtrl) is feasible for macrocells, its degree of isolation is limited in dense, unplanned femtocell deployments.
- Having contiguity in sub-carriers reduces the vulnerability to frequency offsets that can cause interference across cells.
- Isolating resources for data payload and retaining contiguity in sub-carriers to address interference holds promise even in the absence of synchronization.

These inferences lead us to the following set of guidelines that we believe are useful for designing efficient resource management solutions for femtocells. When interference between femtocells does not prevent decoding with lower MCSs, it is better to operate using:

- distributed sub-carrier grouping (e.g., FUSC) to form sub-channels since this reduces the control signaling overhead, while making more room for data transmissions.

However, in the presence of strong interference, it is preferable to operate a femtocell network that

- isolates resources between interfering cells in the frequency domain to increase capacity.
- uses contiguity in sub-carriers (e.g. PUSC) to alleviate the adverse impact of frequency offsets across cells.

VIII. CONCLUSIONS

In this paper, we experimentally characterize interference in OFDMA femtocell networks. Using programmable WiMAX femtocells and commercial clients, we perform measurements with a multitude of design choices including resource isolation, resource mapping and synchronization among BSs. We believe that, this first experimental study on interference characterization in OFDMA femtocells will enable practitioners to design efficient resource management solutions.

REFERENCES

- [1] R. Van Nee and R. Prasad, "OFDM for Wireless Multimedia Communications," *Artech House*, 2000.
- [2] T. Quek, Z. Lei, and S. Sun, "Adaptive Interference Coordination in Multi-cell OFDMA Systems," in *IEEE PIMRC*, 2009.
- [3] R. Chang, Z. Tao, J. Zhang, and C.-C. J. Kuo, "A Graph Approach to Dynamic Fractional Frequency Reuse (FFR) in Multi-cell OFDMA Networks," in *Proc. of IEEE ICC*, 2009.
- [4] J. Yun and K. Shin, "CTRL: A Self-Organizing Femtocell Management Architecture for Co-Channel Deployment," in *ACM MobiCom*, 2010.
- [5] K. Sundaresan and S. Rangarajan, "Efficient Resource Management in OFDMA Femto Cells," in *Proc. of ACM MobiHoc*, May 2009.
- [6] 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.
- [7] TeraSync, "http://www.terasync.net," .
- [8] IEEE 802.16e 2005 Part 16, "Air Interface for Fixed and Mobile Broadband Wireless Access Systems," *IEEE 802.16e standard*.