

A Distributed Resource Management Framework for Interference Mitigation in OFDMA Femtocell Networks

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

Mustafa Y. Arslan
University of California,
Riverside, CA
marslan@cs.ucr.edu

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

Srikanth V.
Krishnamurthy
University of California,
Riverside, CA
krish@cs.ucr.edu

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

ABSTRACT

Next generation wireless networks (i.e., WiMAX, LTE) provide higher bandwidth and spectrum efficiency leveraging smaller (femto) cells with orthogonal frequency division multiple access (OFDMA). The uncoordinated, dense deployments of femtocells however, pose several unique challenges relating to interference and resource management in these networks. Towards addressing these challenges, we propose RADION, a *distributed* resource management framework that effectively manages interference across femtocells. RADION's core building blocks enable femtocells to opportunistically find the available resources in a completely distributed and efficient manner. Further, RADION's modular nature paves the way for different resource management solutions to be incorporated in the framework. We implement RADION on a real WiMAX femtocell testbed deployed in a typical indoor setting. We extensively evaluate two solutions integrated with RADION, both via prototype implementation and simulations and quantify their performance in terms of quick and efficient self-organization.

Categories and Subject Descriptors

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

General Terms

Design, Experimentation, Management, Performance

Keywords

Resource Management, Distributed, OFDMA, Femtocell

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

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.

MobiHoc'12, June 11–14, 2012, Hilton Head Island, SC, USA.
Copyright 2012 ACM 978-1-4503-1281-3/12/06 ...\$10.00.

1. INTRODUCTION

Recent industry reports predict that mobile data usage will grow by a factor of 26 by 2015 with around 40% of the mobile data originating indoors (i.e., homes and enterprises) [1]. The consequent demand for higher capacity has led to smaller cells, called femtocells, that operate using the same OFDMA-based technology as the macro cells [2]. Femtocells offer improved coverage and a cheap and energy-efficient deployment option for indoor usage [3, 4]. They provide numerous advantages to both customers and service providers. For example, customer equipment (e.g., 4G-capable smart phone, laptop) can save energy on the uplink (while enjoying high throughput), since it does not have to transmit to a distant macro base station. For service providers, femtocells offer (i) a cost-effective way of providing coverage, (ii) increased system capacity via spatial reuse and (iii) reduced operational expenses and subscriber churn. Since users can seamlessly migrate between femtocells and macrocells on a licensed spectrum, femtocells offer an attractive alternative to the unlicensed and often congested WiFi systems.

Unlike the planned deployment of macrocell base stations by the cellular operators, femtocells are deployed in an unplanned fashion within buildings. Hence, as femtocell deployments continue to grow, interference will inevitably be a performance limiting factor in a manner similar to that experienced by residential WiFi networks (e.g. [5]). There are two potential sources of such interference for femtocells: (i) interference between femtocells and the macrocell, and (ii) interference between femtocells. This paper focuses on the problem of interference mitigation between multiple femtocells (in range) that stems from their unplanned deployment¹.

Inapplicability of solutions used in macrocells or in WiFi networks: The problem of interference mitigation through efficient resource management is not new and has been extensively studied in different settings. For instance, this problem arises even in the design of macrocell topologies by cellular operators. However, the planned nature of such macrocell deployments leads to very efficient frequency planning and other coordination-based solutions (e.g., fractional frequency re-use), that cannot be applied to our unplanned femtocell scenarios. Efficient design strategies of WiFi networks have also needed to address the problem for unlicensed spectrum. The basic toolkit for WiFi protocols to address interference problems include random access channel contention strategies

¹Other work has focused on the macro-femto interference problem [6], and is beyond the scope of this paper.

that *de-synchronize* competing transmitters in time through methods of carrier sensing and back-off. Unfortunately, femtocells cannot leverage such strategies since their MAC-PHY protocols are required to follow the same *synchronous* channel access methods as their macrocell counterparts. More specifically, standard OFDMA femtocells do not employ carrier sensing based deferral. Hence, they cannot sense the spectrum occupied by other cells to tune themselves onto orthogonal frequencies to resolve interference. Finally, unlike in WiFi, a single OFDMA frame carries data for multiple clients, and our resource management strategies are constrained by these requirements.

RADION – Distributed femtocell resource management: We propose RADION, a *framework* for distributed management of time-frequency resources in OFDMA-based femtocell networks. This framework is designed for the scenario where nearby femtocells cannot explicitly coordinate or interact with each other, and hence is suitable for unplanned residential deployments. Our specific solution requires each femto base station (BS) to intelligently probe availability of resources and use them in an opportunistic and distributed manner. We contrast this design to a centralized resource management solution, called FERMI, that was recently proposed for femtocells [7]. More specifically, FERMI focused on femtocell deployments in enterprise settings, where multiple femtocells were assumed to *cooperate* towards a network-wide objective. In particular, the design in FERMI required the centralized collection of the global network view, constructed through explicit information exchange between the various femtocells. Such cooperation is not realistic for residential settings. Hence the solution in RADION is stylized such that each residential femto BS will try to optimize for its own local objective, without explicitly exchanging information with neighboring femtocell BSs.

Several challenges arise in designing a distributed resource management framework for OFDMA femtocells: (i) OFDMA schedules multiple clients in the same frame. The clients experience different levels of interference and hence, a resource management framework has to account for the characteristics of each client. Specifically, clients with strong interference need to operate on orthogonal resources (i.e., frequency isolation), while clients with weak interference can operate on all frequency resources (i.e., reuse). For the efficient use of resources in OFDMA frames, we need to first differentiate the clients. (ii) Given that multiple clients share frame resources in OFDMA networks, the next challenge is “how to accommodate multiple clients of various classes in the same frame?”. The frame structure has to be carefully managed for various clients considering their interference levels and demands. In particular, each cell needs to determine how much resources can be reused without causing interference to the neighboring cells and how much frequency resources need to be isolated to mitigate the interference from other cells. The frame structure can impact the network wide resource reuse as well, where multiple contention domains are involved. (iii) The resources allocated to one femtocell directly impact the resources for the interfering cells. We need to determine the time and frequency resources of operation for the clients in each cell while accounting for the resources used by neighboring cells. Resource allocation for each femtocell should be adaptive to the network changes, but without explicit coordination due to lack of any central component in the system.

In RADION, we address the above challenges through the specific design of the following three building blocks:

- **Client Categorization:** Active probing is used by the BS to categorize clients into two classes, those that require resource isolation to mitigate the interference and those that can reuse the spectrum, based on their interference levels. Client cat-




Figure 1: Illustration of the WiMAX frame structure.

egorization is a critical module for efficient use of frame resources given that multiple clients have to share the resources in each OFDMA frame.

- **Resource Decoupling:** Once the clients are categorized, RADION is capable of scheduling multiple clients in the same frame by employing a three-zone frame structure. A zone is a dedicated area of the frame in which data to clients can be scheduled. This new frame structure enables frame resource reuse, and hence promotes more efficient use of spectrum. It also decouples resources across contention domains to promote better network-wide reuse.
- **Two-phase Adaptation and Allocation:** The resources used by each BS are determined by a fast and iterative mechanism that is completely distributed. RADION employs a two-phase adaptation for determining jointly the time and frequency resources for each femtocell. Further, it allows for individual execution at each BS without coordination.

Contributions: We design and implement RADION on a WiMAX testbed, consisting of three femtocells in an unplanned indoor environment. To the best of our knowledge, this is the first implementation based design and evaluation of a self-organizing, distributed framework for OFDMA femtocells. We evaluate the various components of RADION to highlight the functionality and accuracy of each. RADION maintains standards compatibility; it is immediately deployable with commercial OFDMA clients (e.g. WiMAX clients). Further, RADION’s components and modules are applicable to the other OFDMA technologies, LTE and LTE-A. We wish to point out that RADION is modular. In particular, depending on the context and the objectives, different resource allocation solutions can be easily incorporated within the RADION framework.

The rest of the paper is organized as follows. In §2, we provide background on WiMAX and OFDMA and discuss related work. §3 discusses the challenges in managing resources in OFDMA femtocell networks and briefly describes how RADION addresses them. In §4, the functional blocks in RADION are described. We evaluate RADION using our testbed and simulations in §5. §6 concludes the paper.

2. BACKGROUND AND RELATED WORK

WiMAX Preliminaries: RADION is implemented on a WiMAX femtocell testbed. In WiMAX, the spectrum is divided into multiple sub-carriers and several sub-carriers are grouped to form a sub-channel. There are two modes of grouping. In *distributed grouping*, sub-carriers are picked as per a pre-determined permutation. This allows a user to see uniform gain and interference across different sub-channels. In *contiguous grouping*, a contiguous set of sub-carriers is grouped to form a sub-channel. Distributed grouping, the default mode in 802.16e standard, is considered in our work.

WiMAX has a two-dimensional frame (see Fig.1) that carries data across time (symbols) and frequency (sub-channels). A combination of a symbol and a sub-channel constitutes a *tile*, which is

the basic unit of resource allocation. Data to multiple mobile stations (MSs) are scheduled as rectangular bursts of tiles in a frame and frames are sent out periodically (every 5ms). Mainly, frames consist of two parts; downlink (from BS to the MS) and uplink (from the MS to the BS). Downlink frames contain the preamble, control and data bursts. While the preamble allows a MS to associate with the BS, the control consists of FCH (frame control header) and MAP. The DL-MAP indicates where each burst is placed in the frame, which MS it is sent to, and what modulation and coding scheme (MCS) should be used to decode it.

Related Work: While OFDMA standards (e.g., WiMAX, LTE) have been drafted recently, related research has existed for some time [8]. Efforts that focus on the macro-femtocells interference [9] and the interference to cell-edge users in OFDMA macrocells exist. The *localized* (cell edge) interference and *planned* cell layouts have aided various interference management solutions [10]. These also include fractional frequency partitioning (FFR) approaches [11], where the spectrum is partitioned into pre-determined static sets. Unlike macrocells, femtocells are deployed in an unplanned manner without coordinated operations and are hence, vulnerable to interference. This necessitates novel interference mitigation solutions. There have been recent studies [9, 12] in this direction but are restricted to theory with several simplifying assumptions that restrict their scope and deployment.

Recently in [7], FERMI, a *centralized* resource management solution for enterprise femtocell deployments was proposed. However, RADION is a completely *distributed* solution targeting unplanned deployment settings (such as residences), where cooperation among femtocells (implicitly assumed in [7]) is not realistic. The concepts of zoning and client categorization were also used in FERMI; however, access to global knowledge at a central controller made it easy to design these functions. RADION allows each femtocell to intelligently determine its zones in a conservative manner. Unlike in FERMI, to account for inaccuracies in distributed zone determination, a transition zone is introduced (details later). This in turn makes the client categorization in RADION also different from that in FERMI; clients that access the medium in the transition zone have to be chosen with care. Due to lack of coordination, accurate client categorization is all the more critical for efficient use of resources.

Diverging from the cellular context, recent efforts show the benefits of OFDMA in WiFi by building systems that enable dynamic spectrum fragmentation [13] and adaptive channel width [14]. In WiFi, APs use one among multiple 20 MHz channels and several conventional distributed channel selection algorithms [5] can be used to configure APs on different channels to avoid interference. However, in femtocells, the entire spectral chunk (say 20MHz) is available to all the cells. They need to operate on mutually orthogonal subsets of frame resources (sub-channels and time symbols) to avoid interference. Thus, resource allocation has to adapt to network dynamics (such as traffic, load etc.). Here, we seek to distributively and quickly determine the resources for each cell. Further, resource allocation constitutes just one component in our broader goal of designing a framework for distributed resource management for femtocells.

3. RESOURCE MANAGEMENT CHALLENGES

The problem of resource management is for each femtocell to distributively determine the frame resources (tiles) that it can use to schedule its clients. Since the resource allocation decisions of one cell impact multiple other cells, efficient mechanisms are needed to quickly converge to a network-wide resource allocation. We first




Figure 2: Introducing transition zone.

discuss the challenges in achieving this objective and briefly describe how RADION addresses each of them.

Client Categorization: OFDMA schedules multiple clients in a frame. Different clients may experience different levels of interference from neighboring femtocells; clients subject to strong interference need to operate on orthogonal sub-channels partitioned across cells (i.e., frequency isolation), while clients with weak interference can use the entire spectrum (i.e., reuse) and still tolerate interference via link (rate) adaptation. Differentiating the clients is key in realizing good spectral efficiencies. A client *categorization* is included within RADION to accurately differentiate between such clients (which can reuse the spectrum *reuse clients*) and those that need spectral isolation (*isolation clients*). Since sensing the medium is not possible with standard OFDMA femtocells, RADION uses an intelligent active probing technique to achieve client categorization with high accuracy. Note that the accuracy of categorization has a direct impact on how efficiently the tiles are utilized.

Resource Decoupling among the Clients: Once the clients are categorized, next question is “how to accommodate multiple clients of different categories (*isolation and reuse clients*) in the same frame?”. Frame resources have to be carefully assigned to both the *isolation* and *reuse* clients; else, it leads to under-utilization of network resources. Each frame can be segmented into two zones, the reuse and the isolation zone, to accommodate the two types of clients (BS1 in Fig.2); clients in the reuse zone will receive data encoded on the entire spectrum, while clients in the isolation zone will receive on a subset of frequencies (determined by an allocation algorithm). The static FFR approaches, using pre-determined sizes of reuse and isolation zone across cells, will however not work in femtocell deployments where interference is pervasive (i.e., not localized). Further, the load of a femtocell as well as interference from other BSs also need to be taken into account in adapting the zone sizes of a particular cell. RADION uses a novel *three-zone* frame structure to address these issues as will be described in §4.2; in a nutshell the additional third zone is a *transition zone* that prevents resource coupling across cells (discussed next).

Resource Decoupling across the Femtocells: When two interfering femtocells have different reuse zone sizes (based on their loads), the larger reuse zone will interfere with the isolation zone of the other cell (BS2 in Fig.2). Having a common reuse zone for the two cells is essential to avoid this. However, irrespective of whether the maximum or the minimum of the reuse zone sizes is chosen as this “common” zone size, it is easy to see that there is under-utilization in one of the cells (either part of the reuse zone or part of the isolation zone is not utilized). More importantly, this coupling can propagate across cells resulting in network-wide under-utilization. RADION’s *transition zone* intelligently localizes this resource coupling and prevents such propagation.

Resource Allocation: Each femtocell has to determine its zone sizes and resource usage in a completely distributed manner. RADION uses an iterative, joint time (zone sizes) and frequency (sub-channels in the isolation zone) resource allocation algorithm that converges to efficient allocations and adapts to network dynamics quickly and efficiently.

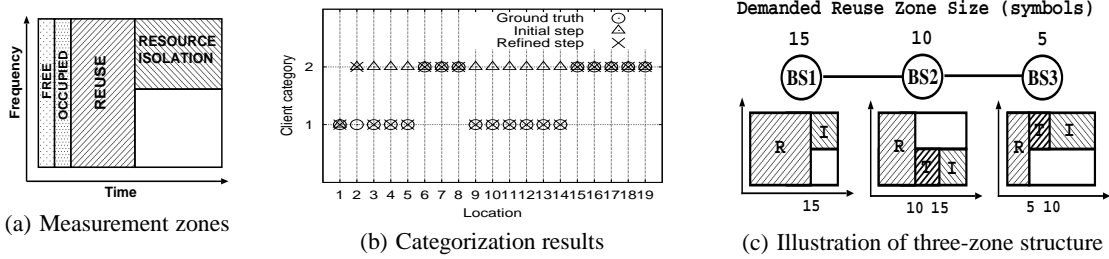


Figure 3: Free and Occupied zones are used for client categorization (a). Refinement step yields accurate categorization results (b).

4. RADION AND ITS COMPONENTS


We next describe RADION in detail by elaborating on the functionalities of its building blocks. A detailed description of our prototype implementation and the testbed used for experiments are given in §5.

4.1 Client Categorization

Initial Step: The first component in RADION categorizes clients into two classes; the first needs just link adaptation (class 1) and the second needs resource isolation (class 2). This is achieved through active probing and Burst Delivery Ratio (BDR) estimation. We define the BDR to be the ratio of successfully delivered bursts to the total number of transmitted bursts. To facilitate client categorization, the frame structure contains two measurement zones of equal size (2 symbols) – *free* and *occupied* zones as depicted in Fig. 3(a). Each BS performs the following two steps towards categorizing a client leveraging these two zones: (i) Schedule the client’s data in the *occupied* zone and schedule it in the *free* zone probabilistically. It is possible that more than one BS would simultaneously schedule its client in the *free* zone. In order to avoid this, a random access mechanism with probability $\frac{1}{n}$ is used where n is the number of interfering BSs. Keep track of the resulting BDR in both zones over K frames. (ii) Determine the normalized throughput per tile in the two zones T_{occ} and T_{free} corresponding to their BDRs. If $T_{free} \geq (1 + \alpha)T_{occ}$ then the client is categorized as class 2; otherwise class 1. Through exhaustive measurements, we set $K = 25$ frames with $\alpha = 0.25$ for high accuracy (>90%) of categorization.

Refinement Step: In the above step, while clients in class 1 are identified accurately, not all clients categorized as class 2 may require resource isolation. With resource isolation, a BS allocates only a subset of resources to a client. For clients with low-moderate interference, link adaptation to cope with interference may be a better option than sacrificing resources through isolation. Thus, to further refine the categorization of class 2 clients, we factor in the loss of resources due to isolation. This was missing in the step (ii), since equal resources were used in the *occupied* and *free* zones. However, the amount of isolated resources available to a cell depends on the resource allocation algorithm. If the resources assigned to the isolation zone is a fraction f of that of the reuse zone, the BS refines the status of a client in class 2 by scheduling the client on resources in the *isolation* zone and determining its normalized per tile throughput in this zone, T_{isol} . It retains the client in class 2 only if $f \cdot T_{isol} \geq (1 + \beta)T_{occ}$; the client is reverted to class 1 otherwise. Here, β (0.05, experimentally obtained) is used to avoid oscillations in categorization. RADION further sub-classifies clients in class 2 as those that benefit significantly from resource isolation (class 2h: $\frac{f \cdot T_{isol}}{T_{occ}} \geq (1 + \alpha)$) and those that benefit marginally from it (class 2l: $(1 + \beta) \leq \frac{f \cdot T_{isol}}{T_{occ}} < (1 + \alpha)$). The benefits of such a sub-classification will be discussed in §4.2.

Evaluation: We consider two cells (1 and 2); clients 1 and 2 belong to the two cells respectively. We generate multiple topolo-



(a) Using two-zone structure (b) Using three-zone structure

Figure 4: Throughput benefits of multi-zone structure.

gies with different levels of interference by varying the location of client 1 in the presence of interfering cell 2. The isolation zone of each cell has the same number of symbols as the occupied zone but operates on an orthogonal half the sub-channels. First, for every location of client 1, its ground truth is determined by scheduling data to the client in the *occupied* and *isolation* zones and determining it as class 2 if $0.5 \cdot T_{isol} \geq T_{occ}$. Then our three step categorization algorithm is executed. The categorization results of the initial and the refinement steps are shown in Fig.3(b). It is seen that the initial step wrongly categorizes clients in ten locations, however, the refinement step corrects most of them. The only erroneous classification (location 2) was due to a change in channel conditions during the process. The clients who need to be in class 2 are correctly classified even with the initial step, while refinement only adds more clients from class 2 to class 1.

4.2 Resource Decoupling

Two-zone Structure: To schedule clients of both classes in the same frame, we use two variable size data zones (*reuse* and *isolation* zones) as in Fig.3(a). The reuse zone operates on all sub-channels and schedules class 1 clients, while the isolation zone schedules class 2 clients on only a contiguous subset of *sub-channels*. To understand the benefits of the two-zone structure, we conduct experiments with two BSs. Each BS has two clients (one in each class), and interferes with the class 2 client of the other BS. The baseline scheme operates the two BSs on two orthogonal sets of sub-channels, with each BS scheduling both its clients within its own subset. This is compared against a two-zone scheme where a BS schedules its class 1 client on all sub-channels, while the isolation zone uses the other half of the frame for its class 2 client on half the sub-channels. We generate various interference topologies and the CDF of the net throughput is plotted for the two schemes in Fig.4(a). We see that with the two zones, class 1 clients can be scheduled in tandem to reuse sub-channel resources effectively, yielding over a 35% throughput gain over the conservative isolation scheme.

Drawbacks of Two-zone Structure: A two-zone structure enables resource reuse, but is insufficient in a multi-cell context. Different cells will have different reuse zone sizes based on the load generated by the class 1 clients. If two interfering cells were to

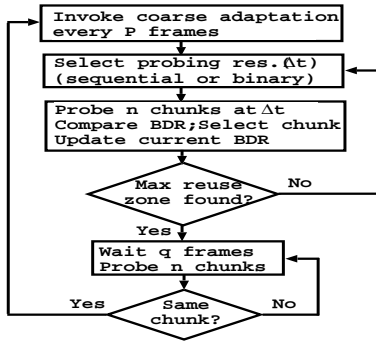


Figure 5: Flow chart of RADION's two-phase adaptation.

operate independently, the larger reuse zone will overlap with the isolation zone of the other cell and hence, interfere with the isolation clients of that cell (Fig.2). Consequently, it is important to use a common reuse zone between interfering cells. For the example topology in Fig.3(c), the result in Fig.4(b) indicates that the throughput of BS2 (with the smaller reuse zone) is degraded when its isolation zone starts right after its reuse zone without accounting for the larger reuse zone of BS1 (2Z curve). This is in comparison to even a simple scheme (2Z-CR curve) that starts the isolation zone of BS2 only after the end of the reuse zone of BS1, leaving the region between the reuse zones of the two BSs unused in BS2.

Need for Common Reuse Zone: RADION uses the maximum of the reuse zone sizes within the interference neighborhood as the common reuse zone for two reasons: (i) since the deployment is uncoordinated and non-cooperative, there is no incentive for a cell to decrease its reuse zone size; (ii) given the absence of sensing by the clients, active probing by the BS to determine resource availability can be employed to *only* determine the reuse zone of a cell with a larger zone size effectively (elaborated in §4.3).

Three-zone Structure: RADION uses a three-zone transmission structure as shown in Fig.3(c). While a BS's *reuse zone* remains the same (for class 1 clients), its *isolation zone* only begins from the end of the common (maximum) reuse zone in its neighborhood. The region between its reuse zone and common reuse zone is the *transition zone*. RADION intelligently picks class 2 clients (using the sub-classification in §4.1) to be scheduled in the transition zone on the same subset of sub-channels as the isolation zone. Specifically, class 2l clients are scheduled in the transition zone, while class 2h clients operate in the isolation zone. There are two benefits to such an approach. First, the transition zone allows selected class 2 clients to reuse resources without incurring significant interference. Second, operating the transition zone on the same subset of sub-channels as the isolation zone prevents the common reuse zone from propagating to the entire network (across multiple contention domains), thereby eliminating under-utilization due to resource coupling.

Evaluation: The benefits of RADION's three-zone structure are evident in Fig.4(b). While both 3Z and 3Z-Opt employ a three-zone structure, 3Z-Opt employs intelligent scheduling of class 2 clients in the transition and isolation zones, 3Z randomly schedules class 2 clients in the two zones. It is clear that the three zone structure with sub-classification yields a 35% throughput improvement.

4.3 Distributed Allocation Framework

We now describe the distributed resource allocation process. The goal of each cell is to determine the size of the reuse (s_r) and transition (s_t) zones as well as the specific contiguous set of sub-channels (C) for operation in the isolation zone. Each BS first classifies its clients into classes 1, 2l and 2h. The preamble used by each BS

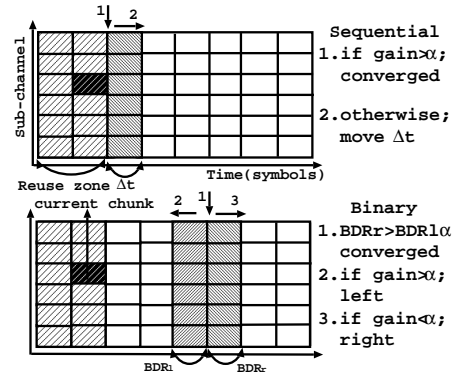


Figure 6: Details of two probing methods.

is chosen from a standardized set (of orthogonal sequences); using this, the clients measure the signal strength to various BSs and hence lock on to a BS with strong signal strength. The same process is used by each client in class 2h to determine its set of strong interferers (received signal strength over a threshold, S_{th}). Using feedback from these clients, the BS determines the super-set of strong interferers. The cardinality of this set including itself (n) determines the fair share allocation of sub-channels ($m = \frac{N}{n}$, where N is total number of sub-channels) for the clients in the isolation zone (class 2h)². Then the BS determines the desired size of its reuse zone (s_r). This is proportional to the relative traffic load from the clients in the two classes. Next, the BS determines the common reuse zone (in time, s_t) in its interference neighborhood as well as the specific set of m sub-channels (in frequency, C) for operation in the isolation zone (details to follow). After determining the resource allocation parameters, clients in classes 1, 2l and 2h are scheduled in the reuse, transition and isolation zones respectively. Determination of the resource allocation parameters in RADION is accomplished with the following joint time-frequency *probe-adapt* mechanism.

4.3.1 Two-phase Adaptation

RADION employs a combination of both coarse and fine time scale adaptations (Fig.5). Every BS picks a period P of coarse adaptation (order of several seconds; thousands of frames). P is picked from a set of large prime numbers to reduce the frequency of overlap of adaptation (and probing) periods across cells. The goal of coarse adaptation is to track coarse network dynamics such as (de)activation of cells/clients, load changes, etc., that happen at the granularity of several seconds. Every P frames, each BS triggers a series of fine time scale adaptations automatically. Once triggered, the goal of fine adaptation is to quickly converge to the right set of parameters (s_t, C), for a given set of network conditions. During fine adaptation, the BS performs a *probe-and-adapt* procedure every q frames till convergence, where q is randomly selected from $[1, 0.1P]$ and operates at the granularity of hundreds of milliseconds. The randomness of q minimizes probing collisions.

Employing coarse adaptation in isolation will result in long recovery times in the event of probing collisions across femtocells. This leads to large periods of degraded performance. On the other hand, employing fine adaptation in isolation will require continuous probing to track network dynamics, thereby resulting in large overhead. RADION strikes a good balance between coarse and fine adaptations; fine adaptation is suspended after quick convergence to an efficient resource allocation and invoked again only in the next

²The solution can be extended to weighted allocations, albeit at the cost of information exchange between femtocells.

Algorithm 1 RADION: Distributed Resource Allocation Framework

```
1:  $m = \frac{N}{n}$ ,  $c_i \in \mathcal{C}$ ,  $s_t \leftarrow s_r$ ,  $b_i$ : BDR,  $e_i$ : counter,  $\forall e_i=0$ 
2: Probe  $c_c$ , update  $b_c$ ,  $s_t += \Delta t$ ,  $f_c$ : current frame
3: while  $(f_c \bmod P) \equiv 0$ 
4:   for  $i = 1 : n$ 
5:     Probe  $c_i$ , update  $b_i$  /*probe  $n$  frequency chunks*/
6:      $b_u : u = \max_{i:c_i \in \mathcal{C}}(b_i)$  /*find the max BDR*/
7:     if  $(b_u > b_c \cdot \alpha) \parallel (b_c > \beta)$ 
8:       Select  $c_f$ : call Algo.2 or 3
9:        $c_c \leftarrow c_f$ , update  $b_c$  (we found  $s_t$ )
10:      Pick  $q \in [1, 0.1P]$ , wait  $q$  frames /*fine adaptation*/
11:      for  $i = 1 : n$ 
12:        Probe  $c_i$ , update  $b_i$ 
13:        Select  $c_f$ : call Algo.2 or 3
14:        if  $c_c \equiv c_f$ 
15:           $\exists c_i \in \mathcal{C} \setminus \{c_f\}$ , s.t.  $b_i > \beta$ ,  $e_i++$ 
16:           $\exists e_i \geq 2$ , pick one of the  $c_i$ ,  $e_i=0$ ; otherwise  $c_i=\emptyset$ 
17:           $c_c \leftarrow c_c \cup c_i$ , goto step 2 /*frequency converged*/
18:        else
19:           $c_c \leftarrow c_f$ , goto step 10 /*re-do fine adaptation*/
20:      else
21:         $s_t += \Delta t$ ,  $c_c \leftarrow c_u$ , update  $b_c$ , goto step 4
```

Algorithm 2 Gibbs Sampler: Frequency Resource Selection

```
1: Temperature parameter:  $T = 0.05$ 
2:  $\forall c_i \in \mathcal{C}$  compute the probability:
    $\pi(c_i) = (e^{-\frac{b_i-1}{T}}) / (\sum_{i=1}^n e^{-\frac{b_i-1}{T}})$ 
3: Sample a random variable  $rand$  with law  $\pi$ 
4: Select  $c_f$  according to  $rand$ 
```

Algorithm 3 Greedy: Frequency Resource Selection

```
1: index:  $i = \arg \max_{i:c_i \in \mathcal{C}}(b_i)$ ,  $c_f \leftarrow c_i$ 
```

coarse adaptation period. Thus, the BS spends a large fraction of its P frames operating on an efficient allocation with its probing and adaptation mechanism constituting only a small portion of it. Further, even during the probing-adaptation procedure, data scheduling to clients is not interrupted and is seamlessly incorporated into it.

4.3.2 Probe and Adapt

The goal of fine adaptation is to ensure quick convergence in the determination of resource parameters in both time (common reuse zone) and frequency (sub-channels in isolation zone) domains. This is achieved with a joint time-frequency adaptation algorithm. Each BS probes a vertical strip of resources (a resource region) in the frame, of size $\Delta t \times N$ (i.e., encompassing all sub-channels in time Δt), where Δt is the granularity of probing in the time domain (few symbols). The frequency domain is further probed in *chunks* of Δf contiguous sub-channels ($\Delta f = m$). Δt and Δf can be varied to tradeoff fine grained allocation and convergence time.

Joint Probing in Time and Frequency: When coarse adaptation is triggered, the BS probes resource regions after its own reuse zone to determine the common reuse zone in its neighborhood (step 2 in Algo.1). The intuition is that since the interfering cell with the largest reuse zone will use all sub-channels till its reuse zone, when frequency chunks are probed within the largest reuse zone, they will exhibit similar (degraded) BDRs, while when probed beyond the largest reuse zone, there will be at least one frequency chunk, whose BDR exceeds those of the other chunks by α (see inference in §4.1). This observation is used by every BS to determine the common reuse zone. Specifically, to probe in a vertical resource region, the BS transmits data to a client in each of n randomly

chosen frequency chunks of size m . Since P varies across cells, and the frequency chunk to be probed is chosen at random, probing conflicts across resource regions are avoided. Each chunk is probed for twenty-five frames and the BDR on each of these chunks is estimated (step 4,5); the maximum BDR (across chunks) is compared to the client's current recorded BDR (step 6,7).

Convergence in Time: We consider two approaches to probing in the time domain: sequential and binary search (Fig.6). In sequential probing (outlined in the pseudo-code), the vertical strip to be probed is advanced sequentially by Δt till a gain exceeding α or a high value ($> \beta=0.8$) of BDR (for BS constituting the maximum reuse zone) is seen compared to the current BDR (time convergence) (step7). Otherwise, the current BDR is updated based on the maximum BDR with the recent probing (step 21). In binary search, two adjacent vertical strips are probed and the BDR with the left and right strips are compared. If $BDR_{right} > BDR_{left} \cdot \alpha$ then size of common reuse zone has been detected (time convergence). Otherwise the maximum value of the BDR (across frequency chunks) is compared with the current BDR to determine if the direction of adaptation should be to the left ($> \alpha$) or right ($\leq \alpha$), and the current BDR is updated only when the region probed is within the maximum reuse zone. If there are multiple clients in class $2h$, they are probed together in each of the chunks and decisions are made with respect to each client. Since different clients may receive interference from different cells, the common reuse zone varies with respect to clients. Time domain probing continues till the common reuse zone for each client in class $2h$ is determined, with the largest common reuse zone determining the termination of the transition zone for the cell.

Convergence in Frequency: Once the common reuse zone is detected, the BS simultaneously has the BDR information on n frequency chunks, with multiple frequency chunks potentially available for operation. We consider two approaches for the selection of a frequency chunk (step 8,9): greedy and Gibbs sampler (Algo.2 and 3). While the greedy scheme is deterministic and picks the chunk yielding the highest BDR, the Gibbs sampler is probabilistic and favors chunks with higher BDR [5]. It has a temperature parameter T , which can be varied with time to provide an annealed version that converges to stable states of low potential (low interference and high BDR). While convergence in the time domain (common reuse zone) can be achieved with high accuracy, frequency domain convergence is sensitive to frequency selectivity and channel errors. Hence, the frequency chunk selected is confirmed by another iteration of fine adaptation, which probes the frequency chunks alone (i.e., no increment of Δt) after q frames (step 10-13). If the same frequency chunk yields the highest BDR, then there is convergence in the frequency domain and the chosen time and frequency parameters are employed for operation till the next coarse adaptation (step 14-17). Otherwise, probing in the frequency domain is repeated every q frames till contention (interference) is alleviated (step 18,19). Thus, by probing vertical strips of frequency chunks, RADION determines both the common reuse zone (s_t) and the set of sub-channels (\mathcal{C} in isolation zone) simultaneously; this leads to quick convergence.

4.3.3 Handling Network Dynamics

Client (dis)-associations impact the traffic load of a cell and consequently, the resource allocations in the isolation zone. Ideally, every cell has to share the frequency resources in the isolation zone of a frame in the contention domains that it belongs to (cliques in the interference/conflict graph), with its ideal share being determined by the size of the largest contention domain that it belongs to. When a new cell is introduced or an existing cell leaves (or traffic ceases)




Figure 7: Deployment and picture of our WiMAX testbed.

the contention domain, the existing share of frequency resources decreases or increases, respectively. However, this change has to be detected by each cell in a completely distributed manner in RADION; this allows cells to contract and expand their sub-channel allocations in the isolation zone. Such a feature is also useful in improving the resource utilization in the network. Note that, since a cell (A) does not have information on its contention domains (requires global knowledge), it computes its fair share (say x , $x \leq$ ideal share) based on its interfering neighbors. However, if one of its neighboring cells (B) belongs to a larger contention domain (hence has a lower share $< x$), then some resources (unused by B) will be under-utilized in cell A's contention domain. The ability to probe and expand resource usage will avoid such under-utilization that is a by-product of distributed operations (smaller granularity of chunk sizes ($< \Delta f$) lowers such under-utilization).

RADION allows a cell to adapt its sub-channel usage as follows. Although a cell selects one of the frequency chunks for operation upon convergence, it keeps track of the BDR in other frequency chunks and hence the potential set of chunks unused by its neighboring cells. It continues to monitor such unused chunks for an additional period of P frames (coarse adaptation period), giving its neighbors enough time to detect and use their fair share. If some of these chunks still continue to be available, then the BS decides to expand its resource usage by adding one of the unused chunks to its allocation (step 16). Adding one chunk at a time, allows other cells in its contention domain to also share the unused resources in a fair manner. This expansion of resource usage will address cases when cells switch off or cease to carry traffic. However, if after expansion, the ceased traffic in a cell restarts or a new cell enters the contention domain, this will be detected in the form of degraded BDRs on the frequency chunks or as a new interferer sensed by its clients. In either case, the BS will contract to its conservative share of sub-channels (in the isolation zone) computed based on its updated set of interfering neighbors and re-run its adaptation algorithm. Any resulting under-utilization in its contention domains will be addressed subsequently through its resource expansion mechanism.

5. SYSTEM EVALUATION

5.1 Testbed and Prototype Implementation

Testbed: Our testbed consists of three femtocells deployed in an indoor environment. The testbed components are shown in Fig.7. We use PicoChip's femto BSs that run WiMAX (802.16e). Our clients are laptops with commercial USB WiMAX cards [15]. All the three cells operate on a 8.75 MHz bandwidth with a carrier frequency of 2.59 GHz. For this frequency, an experimental license has been obtained. We achieve synchronization across femtocells via external GPS modules [16]. Given the cost of procuring programmable WiMAX femtocells, our testbed is restricted to three femtocells. However, we augment our evaluations with simulations, where we study scalability under dense deployments with large number of BSs. Unlike WiFi, femtocells perform syn-

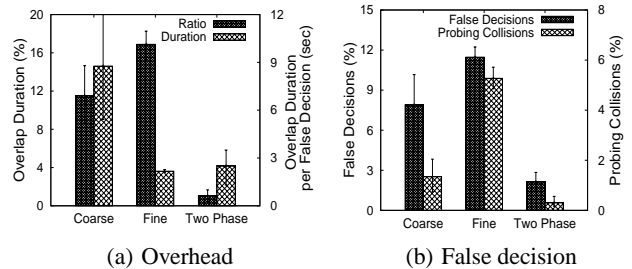


Figure 8: Two-phase adaptation outperforms other schemes.

chronous frame transmissions without carrier sensing. Hence, it makes more sense to generate different interference topologies by varying the locations of clients. This provides a finer control on the inter-BS interference magnitude as opposed to changing the locations of the BSs, and also covers a wide range of scenarios that include both line-of-sight (LOS) and non-LOS links. We also consider ideal link adaptation where we sequentially run the experiment over all MCS levels and record the one delivering the highest throughput for the given topology.

Implementation: RADION is implemented on the PicoChip platform [17], which provides a *base reference* implementation of the 802.16e standard. We significantly extend and modify the MAC scheduler (≈ 2000 lines C code) to realize the components in RADION. Our key modifications include: (1) Client Categorization: We introduce two measurement zones (i.e., free and occupied zones) for client categorization and schedule data bursts in each measurement zone. (2) Resource Decoupling: Each BS tracks its clients' categories to schedule them in the appropriate zone. (3) Probe-and-Adapt Module: The two-phase adaptation algorithm is implemented on each BS; adaptations are triggered at frame boundaries. Further, RADION's modular nature allows us to incorporate the two variants of channel selection as outlined in §4.

5.2 Prototype Evaluations

Having evaluated the first two components of RADION in §4, here we focus on evaluating the two-phase adaptation algorithm. To understand the efficiency of algorithm we first evaluate the adaptation process in time and frequency domains in isolation, followed by their joint evaluation. We conclude the evaluation of RADION's adaptiveness to network dynamics. We create multiple clique topologies where clients in class $2h$ are within the transmission range of other BSs. The BSs operate on frames with 30 sub-channels and 22 time symbols for data. Each experiment is run for at least 10 minutes and is repeated multiple times to generate confidence results.

5.2.1 Frequency Domain Convergence

Here, we fix the reuse zone of all the three cells to be the same, thereby eliminating the need for determining a common reuse zone. Hence, the focus is only on frequency domain convergence in the isolation zone for class $2h$ clients; every cell has to identify and operate on a contiguous set of 10 sub-channels each (3 contending cells). We evaluate RADION's two-phase (coarse + fine) adaptation against the coarse and fine time scale adaptations in isolation. The coarse adaptation period is chosen to be a prime number of frames in $[1000, 6000]$ for each BS and is fixed for subsequent adaptations. Similarly, the period of fine adaptation is chosen randomly in $[1, 600]$. Further, selection of sub-channels after probing follows the greedy approach (Gibbs sampling considered later).

To understand the impact of false decisions on system (throughput) performance, we consider the following metrics: (i) *fractional overlap duration*: the fraction of time that cells operate on over-

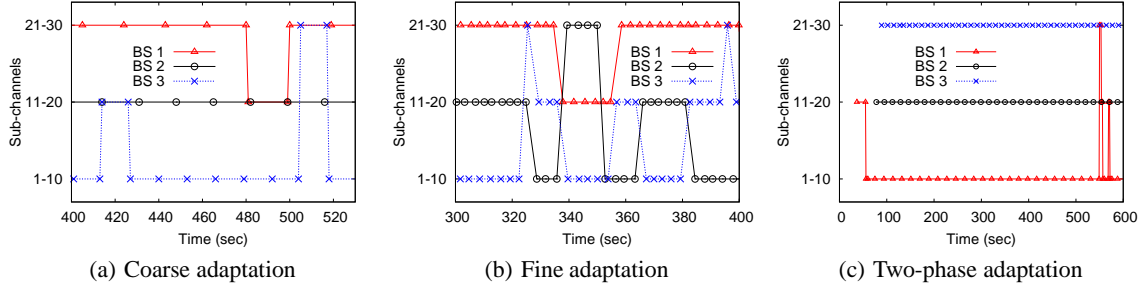


Figure 9: Three convergence patterns.

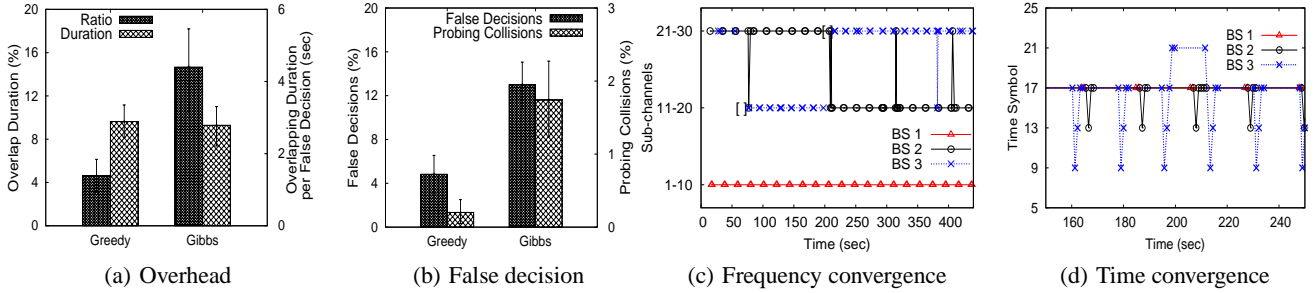


Figure 10: Greedy selection outperforms Gibbs sampler (a and b). Convergence patterns of greedy selection (c and d).

lapping resources (leading to collisions); **(ii) overlap duration per false decision**: ratio of the net duration of resource overlap to the total number of false decisions (resource overlap \equiv collisions); **(iii) fractional false decisions**: ratio of net false decisions to total number of decisions; **(iv) fractional probing collisions**: ratio of net probing collisions to total number of decisions. A probing collision occurs if two or more cells probe at the same time and choose the same resources for isolated operations resulting in a false decision. However, a false decision can also occur due to inaccurate BDR estimates and/or asymmetric interference patterns. We omit the throughput results and rather focus on the convergence to orthogonal resources among the BSs. Throughput degradation only occurs when the BSs operate on the same resource therefore, the *fractional overlap duration* reflects the throughput result indirectly.

The average results of multiple runs of frequency convergence are in Figs.8(a) and 8(b). Coarse adaptation incurs fewer false decisions compared to fine adaptation, but the time to recovery (overlap duration) is large per false decision. In contrast, while the overlap duration per false decision is small for fine adaptation, probing collisions and hence, false decisions increase with the number of active probes. RADION’s two-phase adaptation strikes a good balance between coarse and fine adaptations, resulting only in 2% of the frames colliding and hence outperforms others in all metrics.

To illustrate, we present one particular example of sub-channel convergence as a function of time in Fig.9. We found similar patterns for all other runs. While the cells maintain orthogonality of sub-channels for most part of the experiment, several differences are observed between the schemes. With fine adaptation (Fig.9(b)), at 352 seconds, BS2 switches from sub-channels 21-30 to 1-10, which is already occupied by BS3. This is a false decision due to an inaccurate probing estimate. After BS2 switches to sub-channels 1-10, BS3 switches to sub-channels 11-20 when its next adaptation is invoked (after q frames). Thus, each cell reacts to others’ decisions to avoid interference and maintains orthogonal sub-channel usage. While the number of such switches is small with coarse adaptation (Fig.9(a)), recovery from a false decision is also slow. While recovery is fast with fine adaptation, multiple switches result in both increased probing overhead and false decisions. The

two-phase adaptation (Fig.9(c)) exhibits the best of both adaptations; it incurs fewer switches while also providing fast recovery from false decisions.

5.2.2 Time Domain Convergence

Here, the three cells are pre-assigned orthogonal sets of sub-channels in the isolation zone. We focus on their convergence to the common reuse zone. Recall that RADION probes multiple frequency chunks and compares their BDRs to determine convergence to this common zone. We study both convergence and the impact of the number of probing frames used, for estimating BDR. Sequential search is used for time domain adaptation. The three cells are set different reuse zone demands of 9, 13 and 17 symbols, respectively. BS3, having the maximum reuse zone, will quickly converge to the common reuse zone, while the other two cells will require adaptation. For a given number of probing frames (per chunk), we run the adaptation experiment over 100 times and determine the fraction of cases where the common reuse zone is accurately determined by cells 1 and 2. We repeat this experiment by varying the number of probing frames. We observe that 25 probing frames are sufficient to correctly determine the common reuse zone in over 94.5% of the cases (more frames only provided marginal improvements).

5.2.3 Joint Time-Frequency Convergence

With three BSs (1,2,3) in a clique topology with reuse zone demands set at 17, 13 and 9 symbols respectively, we run the joint time-frequency adaptation process of RADION at each BS. From the results of multiple runs in Fig.10(a) and 10(b), we see that the greedy sub-channel selection yields quick convergence with a low false error rate. The Gibbs sampler incurs a higher fraction of collided frames compared to the greedy. Since both the schemes employ the two-phase adaptation, the average collision duration per false decision is similar with the two schemes (Fig.10(a)). While one might expect the probabilistic nature of Gibbs sampling to yield better convergence [5], this is only true if inferences are based on sensing as opposed to probing; the probabilistic selection increases the number of probing collisions and hence false decisions (Fig.10(b)). Although the greedy approach is deterministic, diver-

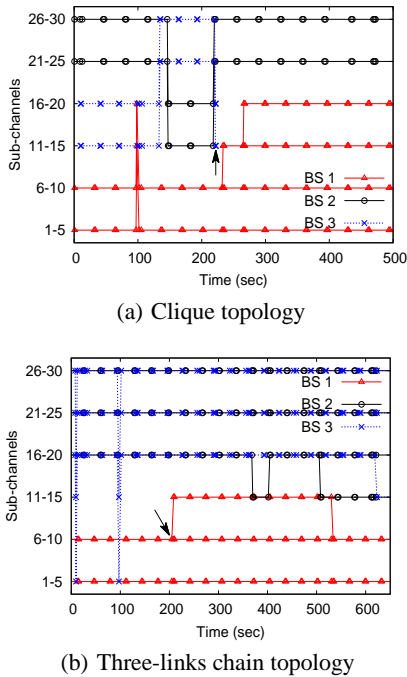


Figure 11: Convergence of RADION with network dynamics.

sity in sub-channel gain across BSs (and their clients) implicitly results in cells picking different frequency resources for their operation. Since both the schemes employ the same time-domain adaptation procedure, their convergence error in the time domain remains similar and less than 5%.

The time-frequency convergence patterns for the greedy scheme are shown in Figs.10(c) and 10(d). Here we present one example of the multiple runs but we confirm very similar patterns for all runs. The frequency adaptations are directly indicated. False convergences in the time domain are indicated by brackets to capture the duration for which the cell operates with a wrong common reuse zone. We see that the number of probing collisions and the resulting quick switch in sub-channel resources to maintain orthogonality in the frequency domain. We omit the pattern of Gibbs sampling for the sake of brevity; however, we highlight that it shows many more frequent switches as the results (Fig.10(a) and 10(b)) show. The time convergence pattern for the greedy approach is expanded in Fig.10(d). We see that BS3 computes the common reuse zone falsely as 21 symbols (instead of 17) at 200 secs; this is corrected in the next coarse adaptation period. Till the reuse zone is corrected, BS3's false decision does not cause interference to BS1's and BS2's class $2h$ clients since it continues to use its isolated resources in its transition zone. However, BS3's class $2h$ clients now incur unfairness as their operational resources are reduced by four symbols. Since the number of false decisions is very small, we let these correct themselves in the next coarse adaptation period.

In summary, we find that the joint time-frequency adaptation process in RADION yields quick and accurate convergence at each femtocell to the common reuse zone as well as to provide orthogonal sub-channels in the isolation zone.

5.2.4 Network Dynamics: Single Contention Domain

We now evaluate RADION's ability to adapt to network changes. First, we consider a single contention domain where all BSs cause interference to each other. In the clique topology of three cells, RADION's adaptation algorithm allows each BS to detect its fair share and converge to orthogonal sets of 10 sub-channels in the isolation

zone. The frequency convergence pattern is shown in Fig.11(a), where frequency resources are probed at the granularity of five sub-channel chunks. RADION's resource expansion mechanism allows each BS to expand its frequency resources in the isolation zone if it probes empty resources for more than a coarse adaptation period. When BS3 is switched off at about 210 seconds (indicated by the arrow), its sub-channels 11-20 become free. BS1 probes the availability of chunks 11-15 and 16-20 but decides to expand its resources only to 1-15, to allow fair access to other cells in the contention domain. However, since the chunk 16-20 remains available in the next adaptation period as well, BS1 continues to expand its resources to 1-20. BS2 fails to grab the chunk 16-20 due to an inaccurate BDR estimate. Thus, while RADION paves the way for a distributed fair sharing of unused resources, distributed operations without information sharing prevents it from controlling the utilization-fairness tradeoff in the best way possible.

5.2.5 Network Dynamics: Multiple Contention Domains

Here we evaluate the adaptiveness in a multiple contention domain by creating chain topology BS3-BS1-BS2; the ideal fair shares of all BSs are 15 sub-channels each. However, since BS1 does not have global information, it computes its fair share to be 10 sub-channels (based on two interfering neighbors) without realizing that its neighbors belong to different contention domains. BS2 and BS3 compute their fair shares as 15 sub-channels each. RADION's resource expansion can help BS1 salvage under-utilized resources. The convergence is shown in Fig.11(b). Initially, while BS2 and BS3 converge to operate on sub-channels from 16-30, BS1 converges to sub-channels 1-10. RADION's resource expansion/contraction feature is enabled after 200 seconds. At this point, BS1 probes the available chunk from 11-15 and expands its allocation to 1-15; this remains stable till 370 secs. At 370 secs, BS2 expands its allocation to 11-30 due to an inaccurate probing, which is immediately rectified in the next adaptation. This again happens at around 510 secs. However, this time, before BS2 rectifies its decision, BS1 responds to interference by contracting its allocation back to its initial fair share of 1-10. This in turn prompts BS3 to probe and expand its allocation to 11-30 (similar to BS2).

The above experiment demonstrates that when a new BS joins the contention domain, it is detected by other cells in the domain; these BSs update (contract) their fair share and run the adaptation process to determine their isolated resources. Stated otherwise, it succinctly captures both the expansion and contraction features incorporated into RADION to track network changes. It also indicates the transition between a fair allocation and one with high utilization. However, to finely control such transitions, information exchange across multiple cells is required, which may not be feasible in residential environments. RADION's *best-effort* utilization-fairness tradeoffs are particularly suited for such environments.

5.3 Evaluation through Simulations

Simulation Set-up: To understand RADION's effectiveness in dense deployments, we resort to simulations. We simulate a single contention domain and increase the number of BSs to stress test RADION's convergence in the frequency domain. We consider two versions of coarse adaptation: (a) a BS picks a fixed prime number of frames P from [8000, 16000], and (b) P is varied (randomly) across adaptation periods. A larger range is chosen for P to avoid excessive probing collisions. For fine adaptation, q is chosen from $[n \times s, 600]$, where n is the number of BSs and s is the number of probes sent on each chunk (set to 10 frames). The variable range is based on the observation that smaller ranges are unfair when a large number of BSs simultaneously probe and adapt. The maximum

	5 BS Max Q: 600 - 1600 - 2600 - 3600	7 BS Max Q: 600 - 1600 - 2600 - 3600	10 BS Max Q: 600 - 1600 - 2600 - 3600	12 BS Max Q: 600 - 1600 - 2600 - 3600
Collision Time Fraction	Greedy: 0.049 - 0.029 - 0.021 - 0.016 Gibbs: 0.131 - 0.094 - 0.073 - 0.062	Greedy: 0.21 - 0.113 - 0.076 - 0.057 Gibbs: 0.309 - 0.248 - 0.204 - 0.17	Greedy: 0.648 - 0.46 - 0.326 - 0.234 Gibbs: 0.619 - 0.524 - 0.46 - 0.403	Greedy: 0.847 - 0.72 - 0.585 - 0.459 Gibbs: 0.759 - 0.685 - 0.615 - 0.567
False Decision Fraction	Greedy: 0.015 - 0.009 - 0.006 - 0.004 Gibbs: 0.04 - 0.03 - 0.023 - 0.019	Greedy: 0.05 - 0.026 - 0.017 - 0.012 Gibbs: 0.075 - 0.058 - 0.048 - 0.039	Greedy: 0.15 - 0.094 - 0.062 - 0.042 Gibbs: 0.129 - 0.102 - 0.085 - 0.073	Greedy: 0.215 - 0.154 - 0.112 - 0.082 Gibbs: 0.157 - 0.128 - 0.108 - 0.095
Collision per False Decision	Greedy: 439 - 657 - 960 - 1217 Gibbs: 429 - 639 - 885 - 1124	Greedy: 408 - 629 - 890 - 1139 Gibbs: 402 - 629 - 857 - 1095	Greedy: 310 - 521 - 751 - 1000 Gibbs: 342 - 547 - 772 - 997	Greedy: 241 - 425 - 632 - 856 Gibbs: 296 - 485 - 691 - 908

Figure 12: Results for the fine adaptation process. Fine adaptation is simulated for different period ranges.




Figure 13: Microscopic results for greedy selection.

range of q is also varied among 600, 1600, 2600 and 3600 frames, to study its impact.

Results: The performance of RADION with an increased number of femtocells is presented in Figs.12, 13 and 14. These results clearly corroborate our findings from the prototype evaluation, thereby demonstrating RADION's scalability in dense deployments. While fewer false decisions are incurred with coarse adaptations (Fig.13(a)), the collision duration per false decision is also longer (Fig.13(b)). The contrary is true with fine adaptation (smaller collision duration per false decision but a higher number of false decisions) as seen in Fig.12. RADION combines the best of coarse and fine adaptation and provides a significant reduction in the time spent in collisions (Fig.14). As with prototype results, the greedy approach outperforms Gibbs sampling; increased probing collisions and hence, false decisions are seen with Gibbs sampling. The table in Fig.12 shows the impact of q 's range on fine adaptation. Increasing the range of q provides more time for estimating BDR on the current frequency chunk. Hence, false decisions due to inaccurate BDR estimates and thus, collision durations are reduced. However, adapting q 's range is a double-edged sword. For small ranges, its performance is dominated by false decisions, while at large ranges it is dominated by collision duration per false decision. Hence, adapting q 's range alone is not sufficient; the two-phase process as in RADION is needed.

6. CONCLUSIONS

We design and implement RADION, arguably the first self-organizing resource management framework for OFDMA femtocell networks. RADION consists of three key building blocks i.e., *client categorization*, *resource decoupling* and *two-phase adaptation and allocation*. RADION allows appropriately chosen clients to opportunistically reuse the spectrum while isolating resources for the other clients in a distributed way. We implement RADION using a WiMAX testbed to show its quick convergence to an efficient resource allocation in real settings. We also demonstrate the scalability and efficacy of RADION in larger scale settings with simulations. We only consider downlink performance, however, a similar approach can be applied to the uplink. As part of future work, we plan to investigate the impact of power control at the BSs.

Acknowledgments

We thank Sampath Rangarajan and Rajesh Mahindra from NEC Labs for their invaluable support and suggestions. J. Yoon and S. Banerjee have been supported in part by the following grants of the

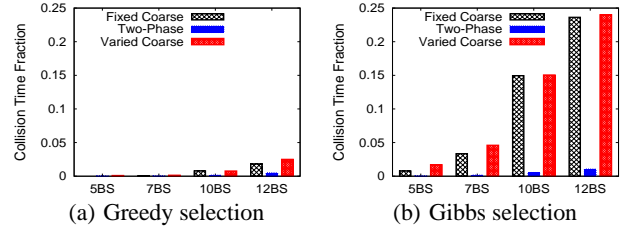


Figure 14: Performance of two-phase adaptation.

US National Science Foundation: CNS-1040648, CNS-0916955, CNS-0855201, CNS-0747177, CNS-1064944, and CNS-1059306. M. Arslan and S. V. Krishnamurthy have been supported by the Multi University Research Initiative grant W911NF-07-1-0318.

7. REFERENCES

- [1] Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2010-2015", Feb 2011.
- [2] R. Van Nee and R. Prasad, "OFDM for Wireless Multimedia Communications", *Artech House*, 2000.
- [3] H. Ekstrom, A. Furuskar, J. Karlsson, M. Meyer, S. Parkvall, J. Torsner, and M. Wahlqvist, "Technical Solutions for the 3G Long Term Evolution", *IEEE Comm. Mag.*, vol. 44, pp. 38-45, 2006.
- [4] S. Yeh, S. Talwar, S. Lee, and H. Kim, "WiMAX Femtocells: A Perspective on Network Architecture, Capacity, and Coverage", *IEEE Comm. Mag.*, vol. 46, pp. 58-65, 2008.
- [5] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot, "Measurement-Based Self Organization of Interfering 802.11 Wireless Access Networks", In *INFOCOM*, 2007.
- [6] V. Chandrasekhar and J. Andrews, "Uplink Capacity and Interference Avoidance for Two-Tier Femtocell Networks", *IEEE Trans. on Wireless Communications*, 2007.
- [7] 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", *MobiCom*, 2011.
- [8] 3GPP, "Technical Specification Group Radio Access Networks; 3G Home NodeB Study Item Technical Report (release 8)", *TR 25.820 V1.0.0 (2007-11)*, Nov 2007.
- [9] J. Yun and K. G. Shin, "CTRL: A Self-Organizing Femtocell Management Architecture for Co-Channel Deployment", In *MobiCom*, 2010.
- [10] D. Lopez-Perez, G. Roche, A. Valcarce, A. Juttner, and J. Zhang, "Interference Avoidance and Dynamic Frequency Planning for WiMAX Femtocells Networks", In *IEEE ICCS*, 2008.
- [11] R. Chang, Z. Tao, J. Zhang, and C. Kuo, "Dynamic Fractional Frequency Reuse (FFR) in Multi-cell OFDMA Networks", In *ICC*, 2009.
- [12] K. Sundaresan and S. Rangarajan, "Efficient Resource Management in OFDMA Femto Cells", In *ACM MOBIHOC*, May 2009.
- [13] L. Yang, W. Hou, Z. Zhang, B. Zhao, and H. Zheng, "Jello: Dynamic Spectrum Sharing in Digital Homes", In *IEEE INFOCOM*, 2010.
- [14] T. Moscibroda, R. Chandra, Y. Wu, S. Sengupta, P. Bahl, and Y. Yuan, "Load-Aware Spectrum Distribution in Wireless LANs", In *ICNP*, 2008.
- [15] Accton, <http://www.accton.com/>.
- [16] TeraSync, <http://www.terasync.net/>.
- [17] PicoChip, <http://www.picochip.com/>.