Coexistence-Aware Scheduling For Wireless System-on-a-Chip Devices

Lei Yang*, Vinod Kone*, Xue Yang†, York Liu†, Ben Y. Zhao* and Haitao Zheng*
*Department of Computer Science, University of California, Santa Barbara, CA 93106
†Intel Labs, Hillsboro, OR 97124

***Abstract*—Today's mobile devices support many wireless technologies to achieve ubiquitous connectivity. Economic and energy constraints, however, are driving the industry to implement multiple technologies into a single radio. This system-on-a-chip architecture leads to competition among networks when devices toggle across different technologies to communicate with multiple networks. In this paper, we study the impact of such network competition using a representative scenario where devices split their time between WiMAX and WiFi connections. We show that competition with WiMAX significantly lowers WiFi's throughput, but this performance degradation is largely unnecessary, and can be attributed to the fact that WiMAX's transmission scheduling does not consider competing networks. We propose PACT, a new coexistence-aware WiMAX scheduling policy that cooperates with WiFi links hosted by its users without compromising its own transmission requirements. We derive PACT's design using an analytical model of network competition, and apply it to design practical WiMAX scheduling algorithms for various traffic classes. We evaluate PACT using OPNET's realistic models for WiFi and WiMAX. Using real network topologies, our experiment results show that PACT significantly improves WiFi performance by up to 17 fold without affecting the WiMAX user experience.**

I. Introduction

Most mobile devices today are equipped with multiple wireless network interfaces, including WiFi, 3G and Bluetooth. Soon WiMAX and cellular LTE will also be added to the list. The extended connectivity opens up new exciting functionality for wireless devices, allowing them to simultaneously connect to different types of wireless networks. For example, a laptop can stream video from the Internet via a WiMAX connection while using its WiFi connection to control wireless peripherals or forward traffic to a local LAN (see Figure 1).

The acceptance of multiple wireless standards has direct implications on hardware design, including how they are implemented into a single mobile device. Building multiple dedicated radios into each device is unattractive for a number of reasons, including cost, increased size and power consumption. Even if implemented, these multiple radios cannot operate simultaneously in practice because of cross-radio interference [1], [2], particularly when their operating frequencies are close. For example, the spectrum mask defined in IEEE 802.11 b/g requires a minimum of 75dB antenna isolation to meet the WiMAX receiver sensitivity requirement [3]. Typical mobile devices, however, only provide 10–30dB isolation [4].

For these reasons, both established chip vendors (Broadcom, Intel) and startups (Wavesat, Altair) have been moving towards the "system-on-chip" architecture, by integrating multiple radio technologies into a single chip [5]. Laptops and handheld

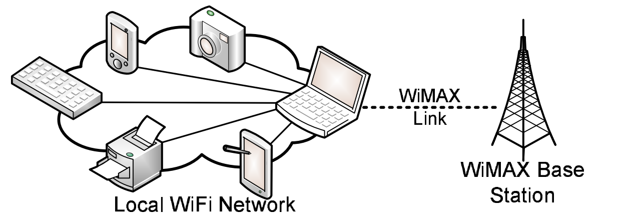

Fig. 1. A laptop splitting time between a WiMAX network for Internet access and its local WiFi network connecting Tablet PCs, PDAs, keyboards, cameras and printers.

devices can switch at will between multiple wireless networks, including WiFi, 3G, WiMAX and LTE. Such an approach lowers manufacturing costs and energy consumption, and further decreases the wireless radio footprint for size-conscious devices.

One important observation is that neither the current nor the new system-on-a-chip solutions allow multiple wireless technologies to operate *simultaneously*. This fundamental limitation has clear implications on performance for both the networks and the user. When connected to multiple networks, a device must split its time across them, creating a new form of local competition between networks. This type of time-multiplexing across networks is already supported by existing wireless standards including IEEE 802.16m [6] for WiMAX and 802.11v [7] for WiFi. These standards allow client devices to coordinate with access points and be temporarily absent from the network without losing existing connections. However, merely allowing active devices to be absent is not sufficient to address the competition between networks. When competing networks' operating protocols or scheduling procedures do not consider their impact to other networks sharing the radio hardware, these networks could have negative interactions and suffer significant performance degradation. Understanding how this competitive environment affects existing wireless technologies is critical to the performance of future devices.

In this paper, we study this problem when multi-radio devices are used to provide concurrent access to both wide-area (WWAN) and local (WLAN) networks. More specifically, we focus our attention on the most immediate and realistic scenario of this type, where laptops using WiMAX for Internet access also use WiFi to host and manage a local network of peripherals and PCs (see Figure 1). These two "overlapping" networks are very different: WiMAX is a centralized, wide-range network with a subscription service model, and WiFi operates for free in local areas providing best-effort services. These differences lead to a significant problem. To retain its WiMAX connection, a mobile device must always listen to the beginning of each WiMAX frame (*5ms* in length) for its transmission schedule. Only if none of its WiMAX downlink or uplink transmissions are scheduled for the current frame, is it allowed to switch back to WiFi mode during the downlink or uplink subframe (*2ms* in length) [1], [2]. Thus, WiMAX user scheduling directly impacts the performance of any competing WiFi networks hosted by these users.

Our experiments confirm that WiMAX scheduling policies can cause significant damage to the throughput of overlapping WiFi networks, making WiFi the big loser in this competitive scenario. This is because existing WiMAX scheduling algorithms can produce usage patterns unfriendly to WiFi connections. They either spread WiMAX transmissions widely across time, leaving little opportunity for WiFi transmissions, or force clients in proximity to operate their WiFi networks in lock-step, thus increasing local interference and contention.

A closer look reveals a surprising result: much of the WiFi performance degradation is unnecessary. In most cases, adjusting WiMAX usage patterns will eliminate the negative impact on WiFi without affecting WiMAX's own performance. In our experiments, making intelligent adjustments can increase WiFi throughput up to a factor of 10. This motivates us to search for new approaches for WiMAX scheduling that meet its own transmission requirements while remaining friendly to overlapping WiFi networks. In developing these approaches, we seek to answer the following questions:

1) *What principles should a WiMAX scheduler follow to meet its QoS requirements and remain WiFi-friendly?*

2) *What information can a WiMAX scheduler collect? How can it be used to positively impact system performance?*

Our Work and Contributions. We first examine these questions by building an analytical model on the WiMAX/WiFi competition. We show analytically that two principles can lead to a coexistence-aware WiMAX scheduling policy. First, WiMAX can "temporally-compress" its user scheduling patterns to increase WiFi's transmission opportunity at no cost to itself. Second, by scheduling radio usage of its users, WiMAX can indirectly limit the level of WiFi contention to improve WiFi fairness and efficiency. Thus instead of competing with WiFi, WiMAX now improves WiFi performance by indirectly managing multiple WiFi networks. Finally, these two principles can be integrated naturally with existing key features of WiMAX user scheduling.

We instantiate our approach as PACT, a practical WiMAX scheduling algorithm that implements the above two principles in a computationally efficient manner. PACT also addresses several practical issues, by supporting multiple WiMAX service classes, adapting to traffic dynamics, and exploiting channel diversity. PACT requires no changes to WiMAX or WiFi devices, and is transparent to existing wireless standards. It operates without any input from overlapping WiFi networks, but can utilize WiFi feedback to further improve performance.

Using both theoretical analysis and experimental simulations, we evaluate PACT's performance under various network configurations. Our analysis shows that PACT has polynomial complexity and yet closely approximates the optimal scheduling solution, the search for which is an NP-hard problem. We also implement PACT on OPNET, which provides both realistic radio propagation models and accurate implementations of WiMAX and WiFi standards. We evaluate PACT using real network topologies from the Google Mountain View and Intel PlaceLab WiFi networks, and make several key findings:

- PACT provides consistently high WiFi throughput while guaranteeing WiMAX QoS requirements, improving throughput by *up to a factor of 17*.
- We integrate PACT with existing WiMAX scheduling algorithms to exploit channel diversity. Under a time-varying WiMAX channel, PACT with diversity improves WiMAX throughput by 30%.
- PACT requires no information from overlapping WiFi networks, but can utilize topology and traffic information to further improve WiFi performance by 10–15%.

Why Change WiMAX Scheduling? PACT addresses network competition by modifying WiMAX base station's scheduling algorithms without introducing additional overhead at clients. It does require commitment of WiMAX operators. An alternative is to keep the base station scheduling intact, but let clients negotiate with their associated networks to define the right schedule. This approach, however, introduces additional overhead at clients. More importantly, it makes WiMAX base stations face difficult challenges of managing many clients with uncoordinated (and often conflicting) absence requests. Since both approaches introduce new complexity at WiMAX base stations, we choose to focus on the first approach to avoid changes to clients. We plan to examine the second approach and compare both in a later study.

II. Related Work

We categorize the related work into *network coexistence* and *WiMAX scheduling*.

Network coexistence. Most works in this area consider network scenarios different from that of PACT: multiple networks compete because their frequency usages overlap and their transmissions interfere with each other, *e.g.* Bluetooth/WiFi. Representative solutions include allocating different frequencies or time slots between networks in a centralized or distributed manner [8], making one technology hop across channels to mitigate the interference [9]. PACT addresses a different coexistence problem: networks operate on different frequencies but their clients toggle between networks in time, and one network (WiMAX) has a higher priority than the other (WiFi). In this case, both frequency hopping and allocation no longer apply. New time-sharing solutions in the form of WiMAX scheduling are required to address the coexistence.

PACT's coexistence scenario was originally discussed in [1], [2]. The work in [2] compared several WiMAX usage patterns in their impact on the overlapping WiFi performance, but did not develop any scheduling algorithm. In [1], WiFi devices with WiMAX connections were compensated by prioritizing their WiFi transmissions over WiFi-only devices. This solution redistributes the negative impact of WiMAX to other WiFi devices but does not address the fundamental WiFi performance loss. PACT is motivated by these prior works, but makes a new contribution by developing effective WiMAX scheduling algorithms to minimize the negative impact.

PACT's coexistence scenario is enabled by the feature of time-multiplexing across networks, supported by IEEE 802.16m [6] and 802.11v [7]. PACT, on the other hand, focuses on how to design network scheduling algorithms on top of this feature to improve network coexistence performance.

WiMAX scheduling. There are many works on cellular and WiMAX scheduling [10]–[13]. Most of them focus on guaranteeing QoS and fairness [14]–[16] or increasing network capacity [17], [18]. PACT, on the other hand, considers a new objective – minimizing the negative impact on the overlapping WiFi networks. As we will show, PACT focuses on choosing the right number of users in each WiMAX frame and thus can be integrated with existing scheduling algorithms.

PACT differs from prior works on network relays [19]–[21], where a relay device forwards traffic between networks, either using both radios simultaneously or the same MAC protocol. PACT addresses two drastically different access methods that compete for radio hardware in time.

III. The Network Competition Problem

In this section we describe the competition between WiMAX and WiFi, examine its impact on network performance, and explore the origins of such negative impact.

A. Problem Scenario and Assumptions

As shown in Fig. 2, we consider a scenario where a WiMAX base station serves a number of laptops and mobile devices. A subset of these devices is equipped with a "compact radio" [3] that supports both WiMAX and WiFi, but can only operate in one at any time. The switching delay between two radio modes is $O(\mu s)$. Compact radios use WiMAX as backhaul to the Internet while simultaneously using WiFi to connect to local peripherals or to host a local wireless hotspot. Having a significantly larger coverage, the WiMAX network overlaps with multiple WiFi networks. The WiMAX and WiFi transmissions operate in different frequencies, and hence can occur simultaneously if they do not involve the same device.

A WiMAX network operates differently compared to WiFi. WiMAX is a wide-area network using centralized scheduling to provide guaranteed services. The WiMAX base station schedules user transmissions in frames and announces the transmission schedule of both directions at the beginning of each frame. In the most common TDD mode, a frame ($5ms$ in length) is divided into a downlink and an uplink subframe.

Because each WiMAX radio strictly follows its transmission schedule, each compact-radio's operating mode is determined by its WiMAX schedule. To maintain its WiMAX connection, a compact-radio must always switch to the WiMAX mode at the beginning of each frame to decode preamble and obtain its transmission schedule, and only switches back to the WiFi mode at the beginning of a sub-frame when it is not scheduled in the current downlink or uplink subframe.

TABLE I
Simulation Parameters

WiMAX		WiFi	
Central Frequency	2.5GHz	Central Frequency	2.4GHz
Frame Size	5 ms	Mode	802.11b
Bandwidth	20 MHz	RTS/CTS	Disabled
Modulation & Coding	64QAM 3/4	TX Power	5 mW
DL:UL Ratio	29:18	PHY Rate	11Mbps

In this paper we consider a conservative scenario where a device uses the current *uplink* subframe ($2ms$) to communicate in WiFi mode only if it finds that it is not scheduled for WiMAX this period. Thus WiFi performance is affected by the WiMAX uplink scheduling. Using 11Mbps 802.11b, it can communicate up to 2750 bytes, reaching an average rate of $11 \cdot \frac{2}{5} = 4.4$Mbps. Finally, when changing its operating mode, a compact-radio will notify and coordinate with its peer WiFi devices to avoid unnecessary WiFi losses [1], [7].

B. The Negative Impact of Network Competition

With the above network model in mind, we now examine the impact of WiMAX/WiFi competition. Using network simulations, we seek to quantify the impact that WiMAX uplink scheduling policies have on overlapping WiFi networks, and to understand the underlying reasons for these results. Existing WiMAX scheduling algorithms are designed to optimize performance metrics such as throughput and delay, assuming dedicated access to the radio hardware. Using OPNET, we simulate a system consisting of one WiMAX cell and 8 WiFi networks as shown in Fig. 2. Each WiFi network has one compact-radio device, which holds a WiMAX and a WiFi connection simultaneously. Each WiMAX connection supports a delay-sensitive video conferencing session, and each WiFi connection carries backlogged data traffic. Table I summarizes our simulation parameters. By default, all the WiFi radios operate on the same channel. We have also examined cases where the WiFi radios operate on different channels.

Figure 3 plots the aggregated WiFi throughput from the 8 WiFi networks when each compact-radio runs a video conferencing session with the WiMAX base station. We see that as WiMAX traffic load increases, WiFi throughput decreases drastically from its normal value of 4.8Mbps to almost zero. After examining the behaviors of WiMAX and WiFi networks, we found that the performance drop can be attributed to two issues related to WiMAX scheduling: *temporal exhaustion* and *spatial imbalance*. In the following, we describe each observation, and discuss the potential performance improvement to make the WiMAX scheduling coexistence-aware.

Cause 1: Temporal exhaustion. We begin by observing the transmission opportunity each single WiFi link has in the presence of WiMAX. We modify the scenario to assign WiFi links with different WiFi channels so that any lack of transmission

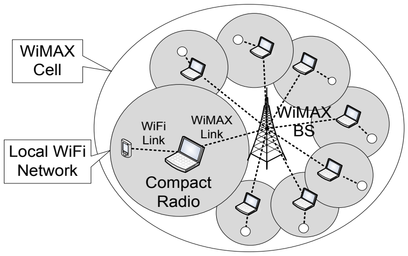

Fig. 2. Competition between WiMAX and WiFi networks that share compact-radios.

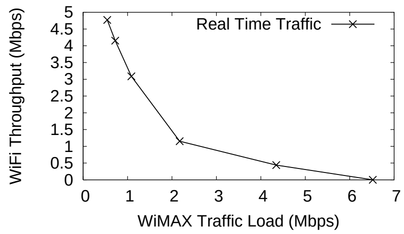

Fig. 3. Aggregate throughput of 8 WiFi networks operating on the same channel. Increasing WiMAX load has a devastating impact on the WiFi throughput.

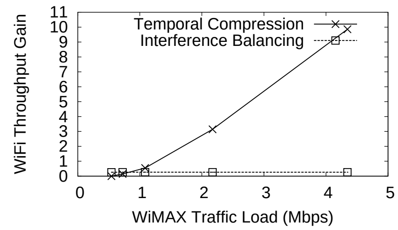

Fig. 4. Temporal compression provides the maximum gain for WiFi, and interference balancing adds a small gain on top of temporal compression.

opportunity is due to WiMAX scheduling. Consider that each WiMAX client requests one uplink transmission every 5ms with a delay bound up to 20ms. A WiMAX scheduler would allocate all 8 clients to transmit one packet in each uplink frame, leaving *no opportunity* for their WiFi transmissions.

In contrast, a coexistence-aware scheduler would allocate each client to transmit 4 packets in every 1 out of 4 frames, allowing the client to operate on WiFi every 3 out of 4 frames. WiFi performance is largely restored without negatively affecting the WiMAX transmission. We refer this new method as "WiMAX temporal compression."

Cause 2: Spatial imbalance. We now turn our attention to the effect of mutual interference between WiFi networks under the impact of WiMAX scheduling. We place all 8 WiFi links on the same channel. Assuming a WiMAX sending rate of 100 packets/second (= 1 packet per 2 WiMAX frames) and an acceptable delay of 10ms, a conventional WiMAX scheduler could allocate 1 WiMAX client in frame #1 and 7 clients in frame #2. This results in 7 compact-radios switching to the WiFi mode in frame #1, creating *heavy mutual interference* and drastically reducing the aggregated WiFi performance.

On the other hand, a coexistence-aware scheduler could schedule 4 compact radios for WiMAX in each of the 2 frames. Now in each frame only 4 compact radios operate in the WiFi mode, thus balancing the interference among WiFi networks over time and improving the overall WiFi throughput.

The Need for Coexistence-Awareness. Clearly, the WiMAX scheduling policy plays a critical role in controlling (1) the transmission opportunity and (2) the contention level of the overlapping WiFi networks. It must be carefully modified to minimize the negative impact on the co-existing WiFi networks. In light of these observations, we compute the possible gain in the aggregated WiFi throughput. First, keeping the WiFi links on separate channels, we compute the gain as the ratio of WiFi throughput solely by adding the temporal compression to the original scheduling algorithm. From Figure 4, the WiFi throughput grows by as much as **11 times**. Next, putting all the WiFi links on the same channel, we quantify the *additional* gain provided by the interference-aware scheduling on top of the temporal compression. Figure 4 shows a relatively smaller gain of 20%. In the next sections, we study both methods and their gains in depth and propose a practical scheduling algorithm to achieve these gains.

IV. Analyzing Network Competition

In this section, we build a formal analytical model on the problem of network competition, and use it to identify the optimal WiMAX scheduling policy that provides WiMAX clients with guaranteed services while maximizing the throughput of overlapping WiFi networks.

To make the analysis tractable, we make the following assumptions. All WiMAX clients are compact-radios that are randomly distributed in a region and experience independent WiMAX traffic. When operating in the WiFi mode, each compact-radio connects with a normal WiFi device, uses the same WiFi channel, and has backlogged WiFi traffic. Focusing on characterizing high-level competition among the two networks, the analysis uses an abstract physical model. It assumes that each compact-radio experiences the same (average) transmission quality, and ignores the potential channel diversity. In Section V we show that the proposed solution can exploit channel diversity to improve WiMAX performance.

A. Modeling Network Competition

We model the network competition problem as a constrained optimization problem. Given WiMAX clients' QoS constraints, it configures the WiMAX schedule to maximize the overlapping WiFi throughput. Next, we define the WiMAX QoS requirement and the aggregated WiFi throughput.

WiMAX Scheduling and QoS. Given N clients and a period of M frames, let $\mathbb{A} = \{a_{i,j}\}_{1\le i\le N, 1\le j\le M}$ represent the WiMAX uplink's user schedule, where $a_{i,j} \in [0,1]$ represents the normalized capacity assigned to a client i in frame j. Let C be the total capacity achievable if a client is assigned with the entire frame. Any client i either sends $a_{i,j}C$ bits in frame j, or is not scheduled ($a_{i,j} = 0$).

A client i also has a QoS requirement defined by (D_i, B_i), where D_i is the delay bound in the unit of frames, and B_i is the minimum number of bits carried per frame. In other words, the QoS constraint requires that the base station must schedule client i to transmit $D_i \cdot B_i$ bits every D_i frames. We rewrite the constraint as:

$$\sum_{j=1+kD_i}^{(k+1)D_i} a_{i,j} \cdot C \ge D_i \cdot B_i, \quad \forall k, 0 \le k < \frac{M}{D_i}, \forall i \qquad (1)$$

Let L define the fraction of WiMAX capacity required to support all N clients: $L = \sum_{i=1}^{N} B_i / C, \quad L \in (0, 1]$.

WiFi Throughput. We now estimate the aggregated WiFi throughput. Note that a compact-radio device i operates in WiFi in frame j only if $a_{i,j} = 0$. Let m_j be the number of clients with $a_{i,j} > 0$, and $n_j = N - m_j$ be the number of active WiFi networks. Using network geometry, we model the spatial reuse of WiFi networks and derive the aggregate WiFi throughput in frame j as a function of n_j:

$$U_{wifi}(n_j) = (1 - (1-\theta)^{n_j})\,\frac{R}{\theta} \tag{2}$$

where R is the throughput of any stand-alone WiFi network, $\theta \in (0, 1]$ is a topology related parameter. As the number of WiFi networks increases, the aggregated WiFi throughput increases due to a higher WiFi spatial reuse. The detailed derivation is omitted here due to space limit. For any given M frames, the aggregate WiFi throughput can be estimated as:

$$U^{agg}_{wifi}(\mathbb{A}) = \frac{\sum_{j=1}^{M} U_{wifi}(n_j)}{M} = \frac{R}{M\theta}\sum_{j=1}^{M}(1-(1-\theta)^{n_j}) \tag{3}$$

While using a specific WiFi performance metric, our solution and results from theoretical analysis apply to other metrics as long as they are monotonically increasing and concave to n_j.

We now define the problem of network competition as a constrained optimization problem: finding the optimal scheduling pattern $\mathbb{A}$ that fulfills all the WiMAX QoS requirements and maximizes the aggregated WiFi throughput:

$$\begin{aligned}
\text{Given} \quad & \{B_i, D_i\}_{i=1}^{N} \text{ and } M \text{ frames}\\
\text{Find} \quad & \mathbb{A} = \{a_{i,j}\}_{N\times M}\\
\text{Maximize} \quad & U^{agg}_{wifi}(\mathbb{A}) = \frac{1}{M}\sum_{j=1}^{M} U_{wifi}(n_j)\\
\text{Subject to} \quad & \sum_{i=1}^{N} a_{i,j} \le 1, \;\; \forall 1 \le j \le M \qquad (4)\\
& \sum_{j=1+kD_i}^{(k+1)D_i} a_{i,j}C \ge B_i D_i, \;\; 0 \le k < \frac{M}{D_i} \qquad (5)
\end{aligned}$$

The constraint (4) defines the WiMAX frame capacity and the constraint (5) defines the WiMAX QoS constraint.

B. The Optimal WiMAX Scheduling Policy

Given the above optimization problem, we investigate the optimal WiMAX scheduling pattern, referred to as $\mathbb{A}_{opt}$, and the complexity of finding it. We show that the optimal solution leads to a unique scheduling pattern (Theorem 1), but finding it is a NP-hard problem (Theorem 2).

Definition 1: A scheduling pattern $\mathbb{A}$ is "Flat" if and only if it satisfies the following conditions: 1) each WiMAX client i is scheduled in one out of every D_i frames, 2) the number of WiMAX clients scheduled in each frame is a constant m.

Theorem 1: If an $\mathbb{A}_{flat}$ exists, then it is the optimal scheduling pattern, $\mathbb{A}_{opt} = \mathbb{A}_{flat}$.

Proof: We prove by showing that if an $\mathbb{A}_{flat}$ exists, it leads to a higher WiFi throughput than any other scheduling pattern $\mathbb{A}$: $U_{wifi}(\mathbb{A}_{flat}) > U_{wifi}(\mathbb{A})$. While multiple flat patterns may exist, they all lead to the same WiFi throughput.

For each $\mathbb{A}$ and each frame j, let n_j be the number of active WiFi networks and m_j be the number of active WiMAX clients, and $n_j + m_j = N$. From (2), $U_{wifi}(n)$ is a concave and strictly increasing function of n. Therefore, we have

$$\begin{aligned}
U_{wifi}(\mathbb{A}) &= \frac{1}{M}\sum_{j=1}^{M} U_{wifi}(n_j)\\
&\le U_{wifi}(\frac{1}{M}\sum_{j=1}^{M} n_j) = U_{wifi}(N - \frac{1}{M}\sum_{j=1}^{M} m_j)
\end{aligned}$$

If $\mathbb{A}$ is not "Flat," it must satisfy at least one of the following conditions: (1) the number of active WiMAX clients varies across frames, $\frac{1}{M}\sum_{j=1}^{M} U_{wifi}(n_j) < U_{wifi}\left(\frac{1}{M}\sum_{j=1}^{M} n_j\right)$ because $U(.)$ is concave; or (2) at least one client uses more than one frame, thus $\frac{\sum_{j=1}^{M} m_j}{M} > m$. Under any of these conditions, we have $U_{wifi}(\mathbb{A}) < U_{wifi}(N - m) = U_{wifi}(\mathbb{A}_{flat})$. ∎

Theorem 1 leads to the same intuition as that in Section III-B: the WiMAX base station should allocate each client with a minimum number of frames to give WiFi networks more opportunity, and distribute clients evenly across frames to maintain a proper level of WiFi contention.

Theorem 2: Finding $\mathbb{A}_{flat}$ is NP-hard

Proof: We first prove the feasibility problem is NP-complete using a reduction from the 3-partition problem [22]. Given $3K$ positive numbers w_i, the 3-partition problem is to decide whether it is possible to partition these numbers into K triples that all have the same sum. We now convert this into an instance of our problem: there are $N = 3K$ clients; each client has the same delay requirement $D = K$ but different bandwidth requirement $B_i = w_i/K$, and the overall WiMAX load is $L = \sum_{i=1}^{N} B_i/C = 1$. If $\mathbb{A}_{flat}$ is feasible, then $3K$ clients are also partitioned into K triples and all have the same sum C. Therefore, the 3-partition problem is equivalent to the feasibility problem of finding $\mathbb{A}_{flat}$. Since the feasibility problem is NP-complete, finding $\mathbb{A}_{flat}$ is NP-hard. ∎

Based on Theorem 2, we can further prove that solving the optimization problem (4) is also NP-hard: if an $\mathbb{A}_{flat}$ is feasible for a scheduling setting, $\mathbb{A}_{flat}$ is the only possible optimal solution (Theorem 1). In this case, finding $\mathbb{A}_{opt}$ is equivalent as finding $\mathbb{A}_{flat}$, and the general optimization problem (4) for finding $\mathbb{A}_{opt}$ is NP-hard.

V. PACT: A Practical Scheduling Policy

Given the hardness of the optimization, we seek to design a heuristic algorithm to approach a "Flat" scheduling pattern. Our design is motivated by Theorem 1: a coexistence-friendly WiMAX scheduler should schedule each client i once every D_i frames and place clients evenly across frames. We propose PACT, an effective and computational-efficient WiMAX scheduling policy to minimize the negative impact to WiFi networks at little or no cost to its own performance. We first present the basic concept of PACT and its practical implementation, and then analyze its performance guarantee.

A. Basic Concept

PACT seeks to schedule uplink WiMAX clients evenly across frames, but prioritize clients by their deadlines to meet the delay and bandwidth requirement (D_i, B_i). One key issue is to determine the number of clients scheduled in each frame. PACT uses the following heuristics: *To satisfy the QoS requirement* $\{D_i\}_{i=1}^{N}$*, the expected number of clients in a frame* ($\overline{m}$) *is defined by* $\overline{m} = \sum_{i=1}^{N} \frac{1}{D_i}$. PACT then opportunistically schedules $\overline{m}$ clients in each frame using an earliest deadline first algorithm [23]. When a frame cannot support $\overline{m}$ clients, the unsupported clients are pushed into the next frame. In the following, we describe the detailed PACT design for practical deployment and analyze the performance bound of using such heuristics-based solution.

B. Making it Work in Practice

Addressing Network Dynamics. The basic PACT assumes that N and clients' QoS requirements remain unchanged over time. In practice, however, clients join/leave the network and change their QoS requirements on-the-fly. PACT adapts to network dynamics by adjusting its scheduling parameters. PACT introduces a sliding window of length W and determines the number of WiMAX clients to be scheduled in each WiMAX frame based on the QoS requirement in the past W frames, smoothing out the impact of network dynamics.

Supporting Mixed Service Classes. The scheduler needs to simultaneously support multiple WiMAX service classes with different forms of QoS requirements. We consider three representative WiMAX service classes: (1) UGS class designed for VoIP sessions with static D_i and B_i; (2) rtPS class designed for video conferencing with static D_i but dynamic B_i; and (3) Best Effort class with no explicit QoS requirement.

To support multiple clients with mixed service classes, PACT processes traffic in a strict order defined by their priority. Although the Best Effort class has no delay requirement, we introduce an artificial delay metric to improve instantaneous throughput and average delay.

Supporting Packet Retransmissions. WiMAX has the option of using hybrid ARQ to recover lost packets. PACT also supports hybrid ARQ by handling ACKs and retransmission packets using the same delay requirements defined by the original scheduler. Because ACKs (for downlink packets) are scheduled at precise times (*e.g.* the 3rd frame after the original packet), PACT does not rearrange them. To reduce the impact of ACKs (which increase uplink load), PACT will apply the same "temporal compressing" concept in downlink scheduling which in return naturally packs multiple ACKs onto one uplink subframe, leaving more time for WiFi.

Exploring Additional Network Information. Note that the basic PACT policy focuses on determining the number of clients to be scheduled in each frame rather than arranging these clients within the frame. Using additional network information, PACT can be combined with sophisticated scheduling solutions that intelligently arrange clients in each frame. We consider three improvements as examples:

- **WiFi topology**. With this information, PACT can further balance the WiFi contention across frames, distributing contending WiFi links into different frames.
- **WiFi traffic**. PACT can indirectly schedule WiFi transmission across frames based on their traffic loads, thus better utilizing WiFi resources.
- **WiMAX clients' channel condition**. PACT can arrange users within each frame to explore channel diversity. For example, after determining the number of clients to be scheduled in a frame, PACT prioritizes clients with better channel quality, using existing proposals [17], [18], [24].

Detailed PACT design. We now describe the PACT scheduling algorithm that replaces the conventional WiMAX scheduling algorithms. For each service class k, the scheduler maintains a task queue Q_k to hold pending WiMAX scheduling tasks sorted by their deadlines. It also keeps a scheduling history H_k of length W frames. For each frame i in H_k, let m_{Ei} represent the number of clients expected to be scheduled in the i^{th} most recent frame, and m_{Ai} represent the number of clients actually scheduled in that frame.

At the start of a WiMAX frame, the scheduler selects clients from a service class k by first computing $\overline{m} = \sum_{i=1}^{N} \frac{1}{D_i}$ using the current QoS parameters, and then deriving the expected number of clients in the current frame m_E as

$$m_E = \overline{m} + \sum_{i=1}^{W} m_{Ei} - \sum_{i=1}^{W} m_{Ai} \quad (6)$$

The scheduler chooses m_E such that the number of clients actually scheduled in each W frames equals to the number of clients expected to be scheduled: $\sum_{i \in H_k} m_{Ei} \approx \sum_{i \in H_k} m_{Ai}$. Next, using the earliest deadline first policy, the scheduler allocates m_E clients from the service class k according to Q_k. If not all m_E clients can be accommodated in this frame due to capacity constraints, it fragments a minimum number of requests and places the residue requests for the next frame. It then updates the actual number of clients scheduled m_A in the scheduling history H_k. Procedure 1 lists the pseudo-code.

Procedure 1 PACT: Coexistence-aware WiMAX Scheduling

Initialize task queue Q_k for each service class k
while A new frame starts **do**
  for each Service Class k (ranked by priority) **do**
    Update task queue Q_k
    Calculate $\overline{m} = \sum_{i=1}^{N} \frac{1}{D_i}$
    Calculate m_E according to (6)
    $m_A = 0$
    while $m_A < m_E$ and FrameCapacity > 0 **do**
      Dequeue a new client i from the task queue Q_k with earliest deadline (If tasks from multiple clients have the same deadline, choose one client randomly or according to the additional network information.)
      Allocate client i
      FrameCapacity -= UsedSymbols for client i
      m_A += 1
    end while
    Update scheduling history H_k by recording m_A and m_E
  end for
  Generate MAP and finish the current frame
end while

C. Analytical Performance Bounds

Being a heuristic driven solution, PACT is not only computational-efficient but also performs within provable bound to the optimal solution. Assuming that the network configurations are static, we show that PACT generates the optimal solution with a high probability when the number of clients N is large; and PACT always produces the optimal solution when the total WiMAX load is small.

Theorem 3: When all N compact-radio clients have the same delay requirement D, and independent and identically-distributed bandwidth requirements B_i with mean B, and the WiMAX load is no more than the network capacity, i.e., $L < 1$, then PACT produces a Flat pattern at a high probability $P_N \to 1$ as $N \to \infty$. In particular, if B_i follows the normal distribution, $P_N \approx \Phi(\sqrt{3N/D}(1-L)/L)$, where $\Phi(x)$ is the CDF of the standard normal distribution.

Proof: To achieve a Flat pattern, PACT must allocate $\overline{m} = N/D$ clients in each of the D frames. This allocation only fails if in at least one frame, the randomly selected N/D clients cannot fit the frame, *i.e.* $\sum^{N/D} D{\cdot}B_i > C$. Therefore, we first examine the probability that any one frame with $\sum^{N/D} D \cdot B_i \leq C$. We define $\overline{B} = \sum^{N/D} B_i/(N/D)$ as the average bandwidth of the randomly selected N/D clients in a frame. Based on the law of large numbers [25], for any $\epsilon > 0$,

$$P(lim_{n=(N/D)\to\infty}|\overline{B} - B| < \epsilon) = 1 \quad (7)$$

Let $\epsilon = (C - BN)/2N$ and because $L = NB/C < 1$, we have $BN + \epsilon N < C$. Combining it with (7), we have $P(lim_{n=(N/D)\to\infty} \sum^{N/D} B_i D < C) = 1$.That is, the probability that randomly selected N/D clients can fit in one frame $P(\sum^{N/D} B_i D \leq C)$ approaches 1 as $N \to \infty$. Because D is finite and $D << N$, the probability that D frames can accommodate N/D clients approaches 1 too. Under the normal distribution, the P_N result follows directly from the central limit theorem [25], thus omitted due to space limit. ■

Theorem 4: When the overall WiMAX load is low, i.e. $C > \lceil\overline{m}\rceil \max_i(B_i D_i)$, PACT always produces an optimal WiMAX scheduling $\mathbb{A}_{opt}$.

Proof: From its definition, PACT always generates a scheduling pattern where 1) each client only occupies one frame in every D_i frames, and 2) either $\lfloor\overline{m}\rfloor$ or $\lceil\overline{m}\rceil$ clients are scheduled in each frame. Similar to the proof of Theorem 1, we can prove that such scheduling pattern leads to the optimal WiFi throughput using the monotonically increasing and concave properties of (3). Thus, we only prove that the allocation satisfies the WiMAX QoS requirements. Since the WiMAX load is low, we only need to show that the delay constraint $\{D_i\}_i$ is met. This is done using existing results on the earliest deadline first scheduler [23]: all tasks can be scheduled before deadlines because the schedulability test is satisfied with an average serving rate $\overline{m} = \sum_{i=1}^{N} 1/D_i$. ■

VI. Evaluation

We evaluate PACT using the OPNET network simulation. OPNET has been widely used to evaluate WiMAX networks given the limited deployment and testbed availability [26], [27]. It provides specialized models for both WiMAX and WiFi matching their IEEE standards.

We implement PACT in the OPNET modeler and compare it to the *built-in* WiMAX schedulers under various network configurations. Specifically, we build a custom compact-radio node with both WiMAX and WiFi connections, and modify the MAC/PHY layers to emulate radio toggling between the two. The system contains one WiMAX base station and multiple WiMAX clients, a subset of which are compact radios. Each compact radio also connects to a WiFi client. We consider three representative WiMAX service classes: UGS, rtPS and best effort, described in Section V-B. We use CBR UDP traffics to simulate these service classes, and set their delay bounds to 20ms (UGS and rtPS) and 100ms (best effort).

When placing the compact radios, we use the measured WiFi topology traces from Place Lab (http://www.placelab.org) and Google-WiFi (http://wifi.google.com) as well as synthetic topologies. The WiMAX base station has a large range of 2km, thus we produce 20 network configurations by placing the base station at different geographic locations. All three types of topologies lead to similar conclusions, thus we only show the Place Lab results unless otherwise mentioned.

A. Addressing Network Competition: Effectiveness and Cost

To understand PACT's effectiveness in addressing network competition, we measure the aggregate WiFi throughput, WiMAX throughput and delay in a network with 8 compact-radios and different WiMAX service classes. Results in Fig. 5 show that PACT increases the WiFi throughput for all three classes, without affecting WiMAX user experience. The WiFi improvements are significant: 1700% for best effort, 200% for rtPS and 26% for UGS, which verifies the advantages of PACT's coexistence-aware scheduling policy.

The results do show that the WiMAX user delay increases. For rtPS and UGS, PACT consistently fulfills the tight delay requirement of 20ms. For best effort, the increased delay could affect users if they use TCP. To examine this impact, we illustrate the WiMAX TCP performance in Fig. 6 by setting different delay bounds. Apparently the performance depends heavily on the choice of delay bound and there is a tradeoff between WiFi and WiMAX throughputs. With a delay bound of 20ms, PACT increases the WiFi throughput by 60% while maintaining a similar WiMAX performance.

An interesting observation is that PACT's gain depends on the WiMAX service class. For best effort and rtPS, the gain is high because their original schedulers tend to schedule as many compact-radios in each WiMAX frame as possible, leaving little opportunity to WiFi. The gain reduces to 26% for UGS because its original scheduler recognizes the constant traffic rate and tends to fit multiple requests from the same compact radio into a single frame. This is similar to PACT's strategy of allocating each client in one out of every D_i frames.

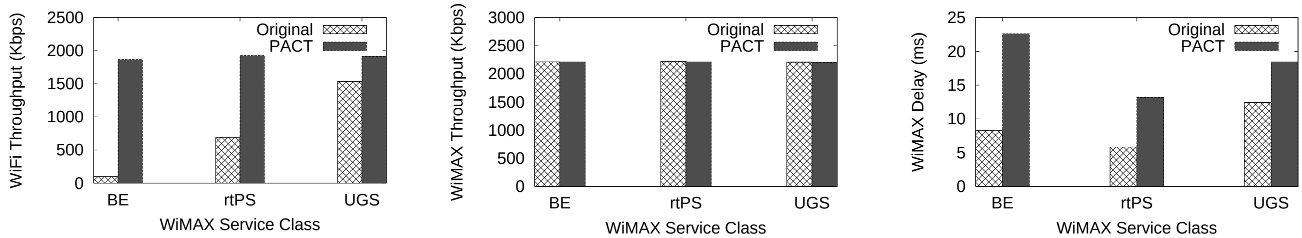

(a) Aggregated WiFi Throughput (b) Average WiMAX Throughput (c) Average WiMAX Delay

Fig. 5. PACT effectively reduces the effect of network competition. It improves the WiFi throughput for best effort by 17 fold and rtPS by 200%, all with no negative impact on WiMAX user performance. The WiMAX delay requirement is 20ms for rtPS and UGS, and 100ms for best effort.

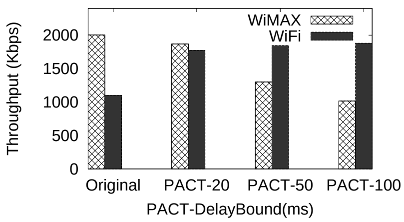

Fig. 6. PACT for best effort WiMAX with TCP. With 20ms delay bound, PACT increases the WiFi throughput by 60% and provides similar WiMAX throughput compared to the original scheduler.

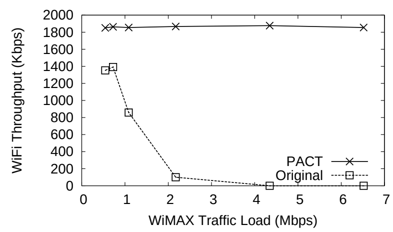

Fig. 7. With best effort WiMAX, PACT provides higher WiFi throughput gain at higher WiMAX loads.

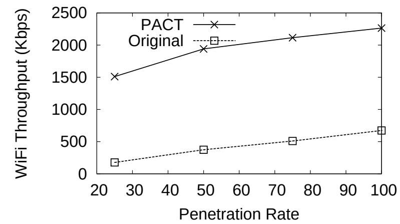

Fig. 8. PACT increases the WiFi throughput with different compact-radio penetration ratios.

Next, we examine PACT under various network settings.

Varying WiMAX Traffic Load and Delay Bound. In Fig. 7, we vary the traffic load of best effort class and compare PACT to the original scheduler. As the WiMAX load increases, the WiFi throughput under the original scheduler reduces drastically due to the temporal exhaustion. Yet PACT leads to a consistent WiFi performance by providing it with adequate opportunity. We also examine the performance by relaxing the delay bound. Results show that PACT further increases the WiFi throughput as the delay bound relaxes, because it now has more flexibility in arranging WiMAX transmissions, leaving more opportunities for WiFi links. We omit this result due to the space limitation.

Varying Compact-radio Penetration Rate. Keeping the total number of clients and the total load in WiMAX constant, we vary the ratio of WiMAX clients with compact-radios. Fig. 8 shows the WiFi throughput for the Place Lab topologies. We see that PACT consistently improves the WiFi throughput, and the absolute improvement quickly converges as the penetration rate increases. This is because PACT consistently achieves the maximum possible WiFi throughput while for the original scheduler the WiFi throughput grows slowly as the penetration rate increases.

Varying WiMAX Traffic Patterns. We also examine PACT in the presence of mixed WiMAX traffic using three traffic combinations: (1) only one service class is present (marked by BE-only, rtPS-only, UGS-only), (2) one of the three service classes is used by 50% of the clients, while the remaining two services are used by 25% clients each. (BE-50, rtPS-50, UGS-50), and (3) all three service classes have equal share of the WiMAX traffic (Equal). From Table II, we see that PACT consistently produces higher WiFi throughput, showing that it supports a wide-range of traffic patterns.

TABLE II
AGGREGATE WIFI THROUGHPUT (KBPS)

Service	Google		Place Lab	
	Original	PACT	Original	PACT
BE-Only	601	2278	673	2264
rtPS-Only	1692	2301	1678	2289
UGS-Only	2122	2284	2107	2275
BE-50	1647	2271	1644	2271
rtPS-50	1809	2287	1797	2268
UGS-50	1782	2283	1773	2270
Equal	1759	2287	1750	2276

B. Enabling WiMAX Channel Diversity

Since OPNET does not report instantaneous channel quality on all subchannels, we perform Matlab simulations to investigate the potential gain by considering the channel quality in PACT's dequeuing process. To make our results consistent with the OPNET results, we use the same modulation and coding schemes and SNR thresholds. We emulate a typical Rayleigh fading environment with a 5ms channel correlation time. Fig. 9 shows that PACT, when combined with channel diversity (PACT-diversity), leads to 30% gain over the basic PACT. We also compare PACT-diversity to an extreme case where WiMAX scheduling utilizes channel diversity to maximize its throughput without considering any QoS and the overlapping WiFi performance. The corresponding throughput performance is an upper bound and PACT-diversity is within 15% gap to this bound, demonstrating its effectiveness.

C. Utilizing WiFi Feedback

Using two case studies, we now examine the potential gain when PACT utilizes additional WiFi information.

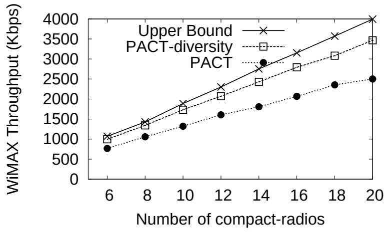

Fig. 9. Integrating PACT with WiMAX channel diversity. If the channel conditions of all WiMAX radios are known at the base station, PACT-diversity can provide 30% gain over PACT.

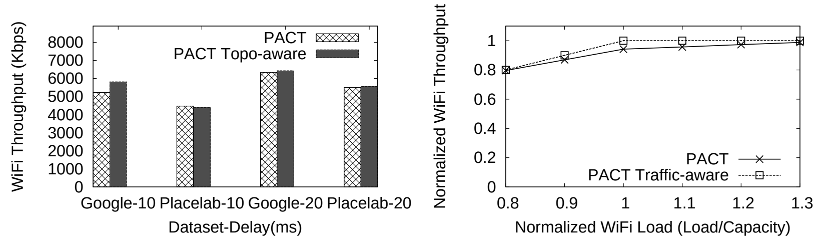

(a) Topology-aware

(b) Traffic-aware

Fig. 10. Utilizing WiFi feedback in PACT. With knowledge about the WiFi topology or traffic information, PACT can further reduce the WiFi contention and improve its throughput. However, the average WiFi throughput gain is marginal (<15%).

WiFi Topology. Given the location of compact-radios, we implement a greedy topology-aware algorithm in the PACT dequeuing process to balance the WiFi interference level. Using UGS traffics, Fig. 10-(a) examines the benefits of considering WiFi topology (described in Section V-B) by comparing it with the basic PACT. The gain is only marginal. This is because the basic PACT scheduler already assigns clients evenly across transmission frames, thus flattening the WiFi interference level.

WiFi Traffic. With the WiFi traffic information, we modify PACT to select active WiFi links in each frame to balance the overall WiFi load. To simulate traffic diversity, we deploy 8 compact-radios where 4 of them have high WiFi traffic demand of H and the rest have low traffic demand of L. We vary the absolute traffic demands but fix the ratio H/L to 10. Fig. 10-(b) shows the normalized WiFi throughput under different loads, where being traffic-aware gets <10% gain. The reason is similar to those for the topology-awareness: the basic PACT scheduler already flattens the WiFi traffic level, leaving little space for improvement.

Overall, the results show that additional WiFi information only leads to marginal improvement under the above scenarios. This implies that PACT's basic scheduler, without requiring any feedback from WiFi, is effective in practice.

VII. Conclusion

We introduce a new problem of network competition caused by compact-radio devices toggling between WiMAX and WiFi networks. We observe that WiFi suffers significant performance degradation, which can be largely prevented by reorganizing WiMAX scheduling patterns. Using this insight, we build an analytical model to identify a coexistence-aware WiMAX scheduling policy that minimizes the negative impact to WiFi. While the problem is NP-hard, we propose PACT, a practical scheduling algorithm that is within provable bounds to the optimum. OPNET simulation results show that PACT improves WiFi performance by up to 17 fold without affecting WiMAX user experience.

References

[1] J. Zhu, X. Yang, and X. Guo, "Disruptive CSMA with credit payback (CP) protocols for multi-radio network," in *Proc. of CrownCom*, 2007.

[2] X. Yang, J. Zhu, and X. Guo, "Study on intermittent wlan consisting of heterogeneous multi-radio devices," in *Proc. of PWC*, 2008.

[3] J. Zhu, A. Waltho, X. Yang, and X. Guo, "Multi-radio coexistence: Challenges and opportunities," in *Proc. of ICCCN*, 2007.

[4] E. Surducan, D. Iancu, V. Surducan, and J. Glossner, "Miniature multiband antennas for hand held WiMAX and WiFi application," in *Proc. of ICEAA*, 2007.

[5] http://www.intel.com/network/connectivity/products/wireless/.

[6] IEEE 802.16, "Part 16: Air interface for broadband wireless access systems," May 2009.

[7] IEEE P802.11v/D3.0, "Part 11: Wireless LAN MAC and PHY specifications, Amendment 8: Wireless Network Management," June 2008.

[8] X. Jing and D. Raychaudhuri, "Spectrum co-existence of IEEE 802.11b and 802.16a networks using the CSCC etiquette protocol," in *Proc. of IEEE DySPAN*, 2005.

[9] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan, "Understanding and mitigating the impact of rf interference on 802.11 networks," in *Proc. of SIGCOMM*, 2007.

[10] N. A. Ali, P. Dhrona, and H. Hassanein, "A performance study of uplink scheduling algorithms in point-to-multipoint WiMAX networks," *Computer Communications*, vol. 32, no. 3, pp. 511–521, 2009.

[11] J. Lakkakorpi *et al.*, "Comparison of different scheduling algorithms for WiMAX base station: Deficit round-robin vs. proportional fair vs. weighted deficit round-robin," in *Proc. WCNC*, 2008.

[12] A. Maltsev and A. Pudeyev, "Multi-user frequency domain scheduling for WiMAX OFDMA," *Mobile Wireless Communications Summit*, 2007.

[13] A. Belghith and L. Nuaymi, "Comparison of WiMAX scheduling algorithms and proposals for the rtPS QoS class," in *Proc. of European Wireless Conference*, 2008.

[14] M. Shreedhar and G. Varghese, "Efficient fair queuing using deficit round-robin," *IEEE/ACM ToN*, vol. 4, no. 3, pp. 375–385, 1996.

[15] A. Parekh and R. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: the single-node case," *IEEE/ACM ToN*, vol. 1, no. 3, pp. 344–357, 1993.

[16] A. Sayenko, O. Alanen, J. Karhula, and T. Hämäläinen, "Ensuring the QoS requirements in 802.16 scheduling," in *Proc. of MSWiM*, 2006.

[17] M. Andrews and L. Zhang, "Scheduling algorithms for multi-carrier wireless data systems," in *Proc. of MobiCom*, 2007.

[18] J. Huang *et al.*, "Joint scheduling and resource allocation in uplink OFDM systems," in *Proc. of ACSSC*, 2007.

[19] H. Luo, R. Ramjee, P. Sinha, L. E. Li, and S. Lu, "UCAN: a unified cellular and ad-hoc network architecture," in *Proc. of MobiCom*, 2003.

[20] S. Deb, V. Mhatre, and V. Ramaiyan, "WiMAX relay networks: opportunistic scheduling to exploit multiuser diversity and frequency selectivity," in *Proc. of MobiCom*, 2008.

[21] K. Sundaresan and S. Rangarajan, "On exploiting diversity and spatial reuse in relay-enabled wireless networks," in *Proc. of MobiHoc*, 2008.

[22] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.

[23] J. Xu and D. Parnas, "Scheduling processes with release times, deadlines, precedence and exclusion relations," *IEEE Transactions on Software Engineering*, vol. 16, no. 3, pp. 360–369, 1990.

[24] J. Huang *et al.*, "Downlink scheduling and resource allocation for OFDM systems," in *Proc. of CISS*, 2006.

[25] G. R. Grimmett and D. R. Stirzaker, *Probability and Random Processes, 2nd Edition*. Clarendon Press, 1992.

[26] S. Ramachandran, C. W. Bostian, and S. F.Midkif, "Performance evaluation of IEEE 802.16 for broadband wireless access," in *Proc. of OPNETWORK*, 2002.

[27] N. Ruangchaijatupon and Y. Ji, "Adaptive scheduling with fairness in IEEE 802.16e networks," in *Proc. of WiCom*, 2007.