

# Symphony: Synchronous Two-Phase Rate and Power Control in 802.11 WLANs

Kishore Ramachandran, Ravi Kokku, Honghai Zhang, and Marco Gruteser

**Abstract**—Adaptive transmit power control in 802.11 wireless LANs (WLANs) on a per-link basis helps increase network capacity and improves battery life of WiFi-enabled mobile devices. However, it faces the following challenges: 1) it can exacerbate receiver-side interference and asymmetric channel access; 2) it can incorrectly lead to lowering the data rate of a link; 3) mobility-induced channel variations at short timescales make detecting and avoiding these problems more complex. Despite substantial prior research, state-of-the-art solutions lack comprehensive techniques to address the above problems. In this paper, we design and implement Symphony, a synchronous two-phase rate and power control system whose agility in adaptation enables us to systematically address the three problems while maximizing the benefits of power control on a per-link basis. We implement Symphony in the Linux MadWifi driver and show that it can be realized on hardware that supports transmit power control with no modifications to the 802.11 MAC, thereby fostering immediate deployability. Our extensive experimental evaluation on a real testbed in an office environment demonstrates that Symphony: 1) enables up to 80% of the clients in three different cells to settle at 50%–94% lower transmit power than a per-cell power control solution; 2) increases network throughput by up to 50% across four realistic deployment scenarios; 3) improves the throughput of asymmetry-affected links by 300%; and 4) opportunistically reduces the transmit power of mobile clients running VOIP calls by up to 97% while only causing a negligible degradation of voice quality.

**Index Terms**—Asymmetry, battery life, interference, mobility, rate adaptation, transmit power control, VOIP, wireless LANs (WLANs).

## I. INTRODUCTION

**P**ROJECTIONS that sales of mobile video phones and WiFi-enabled smartphones will exceed \$100 billion in 2010 suggest that voice and data applications are fast converging from portable onto mobile devices and that WiFi is increasingly being used for last-mile network access [3], [6]. In light of these two trends, WLANs face: 1) higher user mobility; 2) more stringent energy budgets; and 3) higher node densities. While these requirements will likely be addressed through a combination of several mechanisms, we believe that an adaptive transmit power and rate control solution will be an

Manuscript received June 04, 2009; revised November 20, 2009; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor K. Papagiannaki.

K. Ramachandran, R. Kokku, and H. Zhang are with NEC Laboratories America Inc., Princeton, NJ 08540 USA (e-mail: kishore@nec-labs.com; ravik@nec-labs.com; honghai@nec-labs.com).

M. Gruteser is with WINLAB and the Department of Electrical and Computer Engineering, Rutgers University, North Brunswick, NJ 08902 USA (e-mail: gruteser@winlab.rutgers.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2010.2040036

integral component of the overall system. Adaptive transmit power control can compensate for link changes due to mobility, improve spatial reuse by increasing the number of simultaneous transmissions in a given area, and reduce energy consumption for handheld devices, especially when they transmit frequently (e.g., video call from smartphone).<sup>1</sup>

However, despite decades of power and rate control research, few systems exist that jointly adapt transmit power and rate for WLANs on a *per-link* basis (partly due to the lack of hardware support for *fine-grained* power control until recently [35], [47]). In particular, while rate control is performed on a per-link basis, current solutions either use static or coarsely dynamic transmit power configurations on a per-AP or *per-cell* basis [2], [30]. These approaches forego per-link power control's benefits in a mobile environment. Current industry wisdom even seems to view power control as incompatible with VoIP applications [5], further discouraging its adoption. In general, adaptive transmit power control for WLANs faces the following challenges: 1) it can exacerbate receiver-side interference (also known as the hidden-terminal problem) and asymmetric channel access [30], thereby leading to unnecessary packet retransmissions and reduced fairness; 2) it can incorrectly lead to lowering the data rate of a link, thereby increasing the air-time on the channel; 3) mobility-induced channel variations at short timescales make detecting and avoiding these problems more complex.

While some of these problems have been identified previously [22], [24], [30], to our knowledge, no integrated systems solution exists that *simultaneously* addresses all three challenges. Furthermore, no per-link solution exists to address the asymmetric channel access problem in WLANs. The only work on WLAN asymmetry [30] is focused on a per-cell solution that cannot be easily extended to the per-link case, especially in the presence of user mobility, since it converges very slowly (30 s in one of their experiments).

In this paper, we propose Symphony, a novel synchronous two-phase rate and power control system that addresses all the above challenges in a unified framework. The key idea is to use a periodic reference phase during which all nodes operate at maximum power (as if no power control were used) to detect whether power control leads to adverse effects. Symphony is robust to user mobility and can be easily realized on state-of-the-art hardware with no new modifications to the 802.11 MAC protocols. We implement Symphony in the Linux MadWifi driver [29] and demonstrate its efficacy through a

<sup>1</sup>For instance, Broadcom's 802.11b/g interface consumes about 625 and 425 mW when transmitting at 15 and 7 dBm [1], respectively, and 295 mW when receiving. Assuming the device transmits half of the active time and receives in the rest half, this translates to an active-mode battery life increase of up to 27% when operating at 7 dBm instead of 15 dBm.

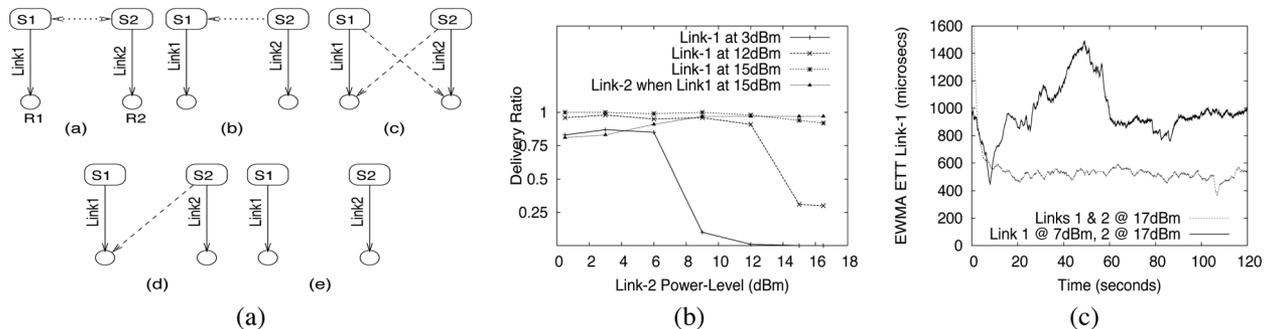


Fig. 1. Problems introduced by power control: (a) Scenarios of interaction between two links. (b) Receiver-side interference on Link-1. (c) Asymmetric channel access on Link-1.

detailed prototype evaluation. Extensive experiments on a real testbed in an office environment and on the indoor ORBIT testbed [41] demonstrate that Symphony: 1) enables up to 80% of the clients in three different cells in an office environment to settle at 50% to 94% (3 to 12 dB) lower transmit power than a per-cell power control solution; 2) opportunistically reduces the transmit power of mobile clients running VOIP calls by up to 97% (15 dB) while causing minimum impact on R-score—a popular performance metric for quantizing voice quality; 3) addresses the problems of hidden terminals and asymmetric channel access while improving the throughput of asymmetry-affected links by 300%; and 4) across four realistic deployment scenarios, increases network throughput by up to 50%.

In summary, we make three contributions in this paper:

- 1) We propose Symphony, the first per-link rate and power control system for WLANs that simultaneously addresses the problems introduced by power control in a comprehensive and easily realizable manner. By implementing Symphony in the Linux MadWifi driver [29] running over off-the-shelf hardware, and requiring no extensions to 802.11 protocols, we demonstrate that the solution is readily deployable.
- 2) As part of Symphony, we propose a novel and simple mechanism based on expected transmission time (ETT) to detect WLAN channel access asymmetry in a distributed manner. Unlike state-of-the-art per-link solutions in the *ad hoc* domain, this mechanism does not require information about the topology [43] or the source of interference [46].
- 3) As part of Symphony, we propose RRAA+, an improved version of the state-of-the-art RRAA rate adaptation algorithm [48], which significantly reduces packet loss and makes links more robust to user mobility. RRAA+ is also useful standalone.

The rest of the paper is organized as follows. In Section II, we describe the challenges of performing adaptive transmit power control in IEEE 802.11 WLANs and discuss related work and their drawbacks. In Section III, we discuss the design and implementation of Symphony, and we present evaluation results in Section IV. Section V discusses open issues and limitations. Finally, Section VI concludes.

## II. PROBLEM FORMULATION

As users increasingly make WLANs their first choice for last-mile network access, both spatial reuse and battery life are crucial metrics to ensure better user experience. With emerging mobile applications leading to increased data transfer over WiFi interfaces, and hardware [7] and protocol [9] improvements reducing the idle-time power consumption of these interfaces, transmit power becomes the dominating factor influencing battery lifetime. Secondly, with increasingly dense deployments of WLANs for continuous coverage to users, mitigating interference to maximize spatial reuse is a crucial design goal. Adaptive transmit power control on a per-link basis promises to improve both the above metrics.

### A. Power Control Challenges

While per-link adaptive power control is beneficial, doing so can be challenging due to several reasons.

*Receiver-Side Interference and Asymmetric Channel Access:* Transmit power control can introduce link asymmetry that leads to two problems: receiver-side interference and asymmetric channel access. Several previous works have already observed these two problems with power control [22], [24], [30], [38]. To provide a more quantitative characterization of: 1) their *effect on performance in realistic settings*; and 2) their *likelihood of occurrence*, consider a canonical network of four nodes—two senders and two corresponding receivers using the same 802.11 channel. Fig. 1(a) identifies the different scenarios of interaction when the two links use different transmit powers. Solid arrows indicate that nodes are in communication range. Dotted arrows indicate that the senders are in carrier-sense range; in scenario (b), S1 can hear S2, but not vice versa. Dashed arrows indicate unintended interference at the receivers. Scenario (a) represents fair channel sharing since S1 and S2 can carrier-sense each other. Scenario (b) represents the case of asymmetric channel access; whenever S2 has data to transmit, S1 does not get a fair chance to transmit. Scenarios (c) and (d) are two instances of receiver-side interference; while the senders are oblivious of each others' presence, packets sent by them collide at their receivers. Finally, scenario (e) is the ideal case of no interference and simultaneous transmissions on each link. Scenarios (b), (c), and (d) for any two links in the network can degrade both link and network throughput as well as fairness.

*Setup:* We use four laptops (with Atheros PCMCIA wireless cards) to emulate this canonical network—two laptops are configured as APs, while the remaining two are configured as clients (one associated with each AP). Both AP–client links are on the same channel (802.11a channel 40). Each AP starts backlogged UDP transfers to its corresponding client, and we fix the 802.11 MAC bit-rate to 54 Mb/s. To emulate receiver-side interference, APs (senders) are positioned such that they do not carrier-sense each other, and clients are placed such that their receptions may be corrupted by interfering transmissions on the other link [scenarios (c) and (d)]. We use packet delivery ratio (PDR), calculated as the ratio of packets successfully received at each client to those sent out by its corresponding AP, as the metric. For asymmetric channel access, we place both receivers such that the network cannot be in scenarios (c) and (d). As the metric, we use the ETT [15], [11] of packets, with the minor modification that we only consider packets that succeed without any retries, to approximate the *channel access delay* for each frame.

*Results:* Fig. 1(b) shows the PDR observed by both links, in the experiment emulating receiver-side interference, when the transmit power of each link is changed. In this figure, the lines “Link-1 at 3 dBm,” “Link-1 at 12 dBm,” and “Link-1 at 15 dBm” show that Link-1 is able to tolerate (due to physical layer capture) varying degrees of interference from Link-2 [scenario (d)] depending on the relative difference in transmit powers. Similarly, the line “Link-2 when Link-1 at 15 dBm” shows that Link-2’s PDR is also affected. In Fig. 1(c), we plot the smoothed (using EWMA) ETT of packets on Link-1. When links are asymmetric [scenario (b)], Link-1 has higher channel access delay, and hence higher ETT, than when both links are symmetric [scenario (a)].

We now address the question, how frequently do the problematic scenarios [(b),(c), and (d)] occur? Since an experimental approach cannot sufficiently answer this question, we take an analytical approach. We derive the probability that each of the scenarios occurs in a random geometrical graph with four nodes as above when the senders employ transmit power control. For brevity, we present the detailed mathematical formulation and analysis in a technical report [39] and just state the results here. Fig. 2(i) shows the probability of occurrence of each scenario with distance between the senders, where the distance is shown as a factor of the communication range of the senders. Fig. 2(ii) shows the sum of probabilities of all problematic scenarios. The graph clearly shows that these scenarios can occur very frequently in a real deployment. Detecting and avoiding these problems in mobile environments is even more challenging since they can be dynamically introduced for short time periods.

*Interaction With Rate Adaptation:* In 802.11a/b/g wireless LANs, senders use one of multiple transmission rates for sending packets. The choice of the rate is determined by an estimate of the channel conditions either through packet loss [23], [27], [29], delivery ratio [48], throughput [11], or signal-to-interference-plus-noise ratio (SINR) estimates [26], [20]. Conceptually, a link is expected to perform well at a chosen rate if the SINR at the receiver is above a threshold (Table I) [26]. Rate selection and transmit power control are tied together; power control without considering rate can reduce the SINR, leading to reduction in rate and, hence, the

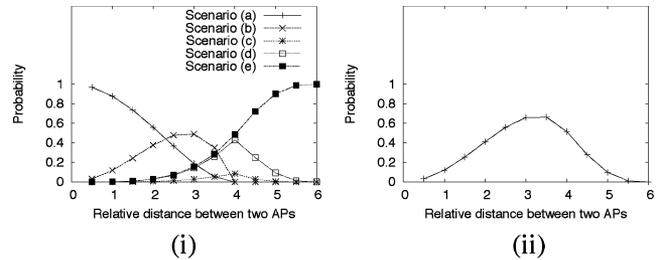


Fig. 2. (i) Probability of scenarios (a)–(e). (ii) Total probability of problematic scenarios [scenarios (b), (c) and (d)].

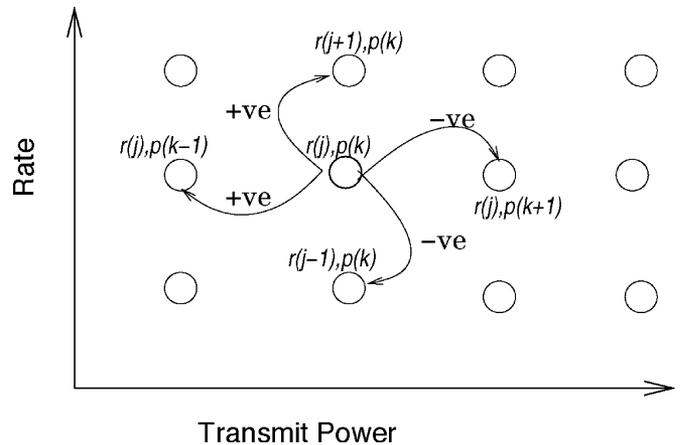


Fig. 3. The rate and power adaptation dilemma.

TABLE I  
SINR VERSUS RATE (Mb/s) FOR BERS  $\leq 10^{-5}$

SINR Range (dB)	Rate	SINR Range (dB)	Rate
$\geq 24.56$	54	[10.79,17.04]	18
[24.05,24.56]	48	[9.03,10.79]	12
[18.8,24.05]	36	[7.78,9.03]	9
[17.04,18.8]	24	[6.02,7.78]	6

link and network throughput. In this paper, we take a systems perspective and choose the (minimum) power level for a link that does not compromise the achievable rate. From the table, it can be seen that for supporting 54 Mb/s, the transmit power can be reduced until it reaches the SINR threshold limit of 24.56 dB. Similarly, say, if 54 and 48 Mb/s cannot be supported even at maximum allowed transmit power, then power can be reduced until the SINR approaches 19 dB, which still allows operation at 36 Mb/s.

While such rate and power selection is easily realizable with precise knowledge of receiver SINR [26], reliable SINR measurements and reports in the presence of mobility cannot be achieved at fine timescales due to their overhead. Consequently, we rely on *estimating* the channel conditions based on the delivery ratio of a window of packets, similar to past works [48]. Such an approach, however, makes rate and power selection nontrivial. To illustrate, consider Fig. 3. If the link is in a state of rate and power allocation  $(r_j, p_k)$  at a given instant, and the delivery ratio deteriorates (negative feedback), the reaction can either be to reduce rate or increase power. While increasing

TABLE II  
TAXONOMY OF EXISTING TRANSMIT POWER CONTROL ALGORITHMS

Domain	Solution	Granularity	Realization	Deployability	Objective		Rate Adaptation	Channel Access Asymm.	Hidden Nodes
					Energy	Capacity			
WLANs	[Sheth02] [44], [45]	Per-link	✓	✓	✓	✓	✓	×	×
	MiSer [37]	Per-link	×	✓	✓	✓	✓	×	✓
	PARF, PERF [8]	Per-link	✓	✓	×	✓	✓	×	×
	[Chevillat05] [12]	Per-link	×	✓	✓	✓	✓	×	×
	[Mhatre07] [30]	Per-cell	✓	✓	×	✓	×	✓	×
Adhoc Networks	PCMA [31]	Per-link	×	×	×	✓	×	×	✓
	BASIC, PCM [22]	Per-link	×	✓	✓	×	×	×	✓
	PCDC [32]	Per-link	×	×	✓	✓	×	×	✓
	POWMAC [33]	Per-link	×	×	✓	✓	×	×	✓
	SHUSH [46]	Per-link	×	×	✓	✓	×	✓	✓
	PRC [26]	Per-link	×	×	×	✓	✓	×	×
	TACP [43]	Per-link	×	×	×	✓	×	✓	✓
	COMPOW [34]	Per-network	✓	✓	×	✓	×	✓	✓

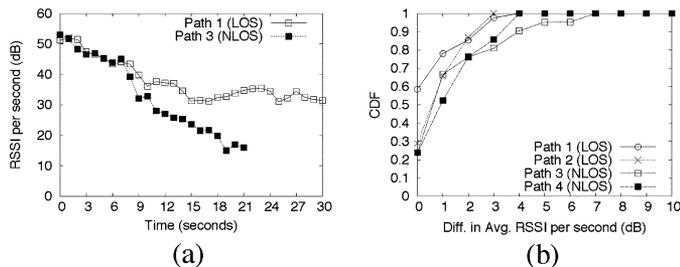


Fig. 4. Effect of mobility (at 0.75 m/s) on RSSI.

power appears to be a natural choice (as in PARF [8]), it is possible that even at the maximum power, the current rate cannot be supported, in which case reducing rate is the right choice. Lack of knowledge of whether a rate can be supported by increasing power to the maximum can increase the convergence time, which is prohibitive in the presence of mobility. A similar dilemma exists for positive feedback.

**Mobility:** Adaptive transmit power control becomes more challenging in the presence of user mobility. Link conditions change frequently due to distance-based path loss, short-lived hidden terminals, and destructive multipath interference at certain locations in a user's path.

To illustrate and quantify these effects, we conducted an experiment with one moving client, at walking speed (approximately 0.75 m/s), sending voice call traffic (50 packets per second) to a stationary AP. The client moves away from the AP along four different paths in our office building, remaining in line-of-sight (LOS) of the AP on two of the paths and in non-line-of-sight (NLOS) on the others. Fig. 4(a) shows the average RSSI per second over time for one LOS and one NLOS case. In Fig. 4(b), we show the CDF of the difference in avg. RSSI per second. Although 95% of the time, the average RSSI in 1 s changes by at most 5 dB, we also observed changes up to 15 dB within 2 s.

In such environments, both rate and power control algorithms need to address the following questions effectively: 1) *How frequently* should the adaptation take place? 2) *At what granularity* should rate and power be adapted? For the first question, solutions have to strike the right balance between reliability and responsiveness: Waiting for enough samples avoids reacting to short-lived drops in link conditions, while waiting too long can also be detrimental to performance. For the second case, using

a coarse granularity allows adapting less frequently, but compromises on battery life and spatial reuse, whereas fine-grained changes require frequent adaptation.

### B. Existing Solutions and Drawbacks

Transmit power control and rate adaptation are well-researched topics in wireless networks. However, our study reveals that no solution provides a comprehensive and easily realizable approach to *simultaneously* address the challenges discussed in the previous section. In particular, our survey across the fields of WLANs, ad hoc, cellular, and sensor networks research, reveals that: 1) there are no existing per-link solutions to address the asymmetric channel access problem in WLANs; and 2) although approaches exist that address a subset of the challenges, Symphony is the first system to simultaneously address the complete set and provide a solution for mobility-enabled WLANs. Table II provides a taxonomy of the relevant related work and identifies their drawbacks. Solutions having deployability constraints due to significant protocol modifications or impractical assumptions, such as the requirement of precise interference measurements, are marked by × under “Deployability.” Solutions not realized in practice and hence not addressing system-level challenges are marked by × under “Realization.” Remaining columns provide information on the granularity of each solution, its objective (reduced energy consumption or increased capacity or both), and whether it addresses receiver-side interference, asymmetric channel access, and jointly adapting rate.

Several per-link solutions in the *ad hoc* domain address channel access asymmetry (Table II). The proposed solutions assume that the hidden/exposed node can be identified so that a message can then be sent to this node. However, the reuse of these solutions in existing WLANs is not possible due to the infeasibility of identifying the source of interference at fine timescales in the presence of mobility (hence marked by × under “Realization”). Among per-cell [30] and per-network [34] solutions, COMPOW [34] addresses the asymmetric channel access problem by making every node in an ad hoc network transmit at a common optimum power. In WLANs, Mhatre *et al.* [30] propose maintaining symmetry by jointly tuning the CCA threshold and transmit powers to maximize network-wide throughput for a *given placement* of APs and clients. However, these solutions do not lend themselves well to be extended for the per-link case for two reasons: 1) they take

a long time to converge after a topology change (e.g., Mhatre *et al.*'s [30] solution takes 30 s in one of their experiments), which is unacceptable in the presence of user mobility; and 2) they operate many links in the network at a significantly higher transmit power than necessary since, in either case, the weakest link determines the transmit power. In general, we note that per-cell approaches have to necessarily degenerate into per-link approaches to efficiently deal with mobility. Finally, approaches proposed for other domains such as sensor networks [21], [28] or CDMA cellular networks [13], [17], [18] cannot be reused mainly because they rely on different multiple access techniques. Specifically, approaches in CDMA cellular networks cannot address channel access asymmetry and require accurate interference estimation at fine timescales, which is infeasible in mobility-enabled WLANs without protocol modifications.

With regards to joint rate and power adaptation for WLANs, to date, we are aware of techniques that address only a subset of the problems and have mostly been implemented in simulators [8], [12], [37]. In contrast, Symphony is the first effort to implement a complete system. The use of transmit power to mitigate the hidden terminal problem [50] is also relatively novel as compared to the use of RTS/CTS messages.

We acknowledge that the related work discussed here is by no means complete and refer the reader to our survey in [39] for more details.

### III. DESIGN AND IMPLEMENTATION

This section describes the design and implementation of Symphony, a synchronous two-phase rate and power control framework to increase the battery life of mobile devices and improve spatial reuse, while addressing several challenges with adaptive transmit power control.

#### A. Overview

Due to the possible adverse effects of power control, the goal of Symphony is to control the transmit power and rate of each link in a WLAN such that *the link's performance is at least as good as in the baseline maximum-power network*. At the core of Symphony is a synchronous two-phase execution (Fig. 5) strategy in which all nodes (APs and clients) in the WLAN cycle through two phases in synchrony—the REFERENCE (REF) phase and the OPERATIONAL (OPT) phase. In the REF phase, Symphony estimates for each link the best achievable performance, and in the OPT phase, it tunes the link to the lowest transmit power to achieve the same performance as in the REF phase. Due to mobility (of users or the environment), the best attainable performance may continuously change, and a power/rate setting may suddenly be affected by asymmetry. The reference phase provides a convenient solution to periodically verify that power and rate control have not unnecessarily degraded system performance.

Similar to most rate control algorithms, Symphony executes on each unidirectional (*sender, receiver*) link at the sender side. The sender can be either an AP or a client, and the receiver a client or an AP. In the REF phase, each sender performs rate adaptation for each link at the maximum power to choose the best data rate for the current channel conditions. In the

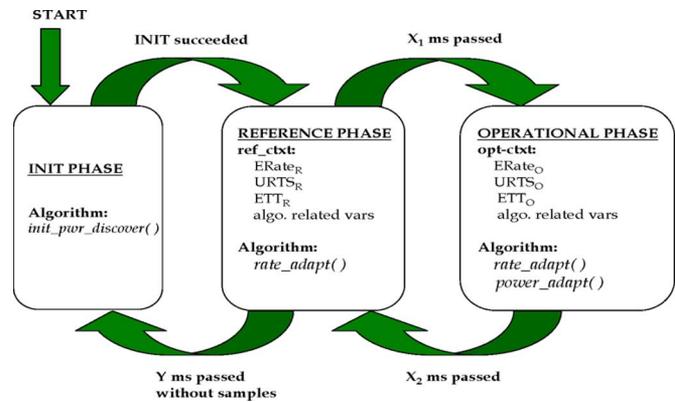


Fig. 5. Two-phase synchronous strategy.

OPT phase, the sender performs both rate and power adaptation. The rate and power adaptation algorithms maintain two contexts—*ref\_ctxt* and *opt\_ctxt*, one for each phase for each link. Each context contains several performance metrics and other variables needed for executing the rate and power adaptation algorithms. We choose three metrics—exponential weighted moving average (EWMA) of data rate, utility of RTS, and EWMA of ETT—to help detect and avoid the problems outlined in Section II. The performance metrics in the *ref\_ctxt* serve as reference values for the OPT phase. In the OPT phase, each link is tuned to the lowest power such that each metric in the *opt\_ctxt* is *no worse* than the corresponding metric in the REF phase by more than a threshold.

Before entering the REF-OPT cycle, each link goes through an INIT phase, in which the sender starts from the minimum power level and rapidly discovers the initial power level necessary for communication with the receiver. To avoid affecting applications, we use probe packets at the highest rate and additionally increase power for each packet till a probe packet succeeds in reaching the receiver. If we reach the maximum power and still do not succeed, we lower the rate and start over at the lowest power. After succeeding, the sender initializes the OPT phase with the successful power level and enters the REF phase with the successful rate at the appropriate synchronized time. If: 1) a packet transmission does not succeed even after 11 retries; and 2) the sender is idle for a threshold number of seconds (= 2 in our prototype), we determine that the current rate and power information is stale (e.g., due to mobility) and reset the sender to the INIT phase.

We achieve synchronized phase execution on all APs and clients in two steps. First, the APs are synchronized to a global real-time clock by a central controller.<sup>2</sup> The controller configures the lengths of the two phases on each AP and specifies at what point in time the phases should start executing. Second, for each phase change, each AP broadcasts a message (at maximum power) informing the change to the clients, and the clients switch phase. These broadcast messages are sent at high priority to ensure minimum skew across nodes; in our prototype, we use the hardware queue reserved for voice traffic in the Atheros cards.

<sup>2</sup>The controller and thin-AP architecture is the most common way WLANs are built today.

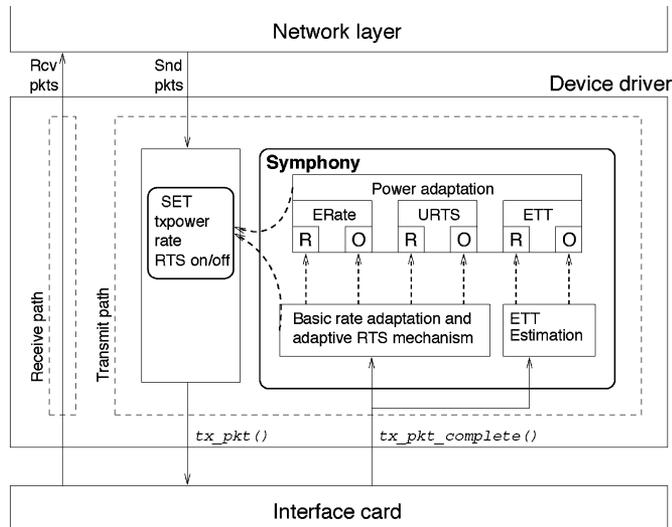


Fig. 6. Architecture of Symphony. The blocks R and O represent REF and OPT contexts.

We implement Symphony in the MadWifi driver 0.9.3.1 [29]. Fig. 6 shows the architecture of Symphony. As shown, Symphony executes in the transmit path. We represent rate adaptation as a separate block to make Symphony extensible. Any rate adaptation algorithm can fit into Symphony as long as it executes in the two contexts and provides rate information for power adaptation. Similarly, different mechanisms can be implemented for ETT estimation and determination of RTS utility. In what follows, we describe the important components of Symphony.

### B. Rate Adaptation

Rate adaptation involves choosing one of several transmission bitrates supported by 802.11a/b/g standards based on the channel conditions. In our context, a rate adaptation algorithm should satisfy at least three requirements: 1) it should be agile to user mobility in typical WLAN environments; 2) it should converge to an appropriate rate for each link rapidly to help with power adaptation; and 3) it should not drop rate due to receiver-side interference, but instead aid the power adaptation algorithm to correct it.

Rate adaptation is a well-researched topic in 802.11 wireless networks [11], [19], [20], [23], [27], [29], [39], [48]. Wong *et al.* [48] and Ramachandran *et al.* [40] provide a survey of several rate adaptation algorithms. In this paper, we analyze two state-of-the-art algorithms in detail: *SampleRate* [11] and *RRAA* [48].

*SampleRate*: The *SampleRate* [11] algorithm selects a transmission rate that minimizes the mean packet transmission time. *SampleRate* maintains an EWMA of expected packet transmission time (ETT) for each rate. For each successfully sent packet, the ETT is updated based on the number of retransmissions, packet length, and protocol timing overheads. Periodically, *SampleRate* attempts transmission at other rates, and if these sample transmissions indicate lower mean transmission time at other rates, the algorithm switches the rate.

*RRAA*: This algorithm [48] uses short-term loss estimation of 802.11 frames (in a window of tens of frames) to opportunistically

guide rate adaptation. The basic RRAA algorithm works with *high* and *low* thresholds  $HT_k$  and  $LT_k$  on loss rate at the current data rate  $k$  selected; Algorithm 1 depicts its behavior. Furthermore, it uses selective RTS/CTS for avoiding unnecessary rate adaptation in response to collision-induced losses (i.e., receiver-side interference). Briefly, selective RTS/CTS works as follows: on detecting loss of frames,<sup>3</sup> RRAA enables RTS/CTS on a selective number of frames. If RTS/CTS reduces the frame loss, data rate is not reduced, instead RTS/CTS is increasingly enabled on a greater number of frames. Otherwise, RRAA determines that losses are not because of receiver-side interference and reduces rate while disabling RTS/CTS. Note that, while turning on RTS/CTS on *all* frames avoids receiver-side interference completely, it also reduces the link and network throughput due to protocol overhead and, hence, is not generally enabled in practice.

### Algorithm 1 RRAA

- 1: **if** ( $loss > HT_k$ ) **then**
- 2:  $k \leftarrow next\_lower\_rate$
- 3: **else if** ( $loss < LT_k$ ) **then**
- 4:  $k \leftarrow next\_higher\_rate$
- 5: **end if**

Basing rate adaptation on loss estimation over a window of packets helps a transmitter to not react adversely to one or a few packet losses, which are more common with mobility and power control. The second feature of not reducing rate in response to collision-induced losses is a useful feature [25], [48] and is missing in *SampleRate*. One drawback of RRAA, however, is that it does not converge to a particular rate if the next higher rate is inappropriate. For instance, if the 54 Mb/s rate causes frame loss higher than the high threshold, and 48 Mb/s causes frame loss lower than the low threshold, RRAA flips between 54 and 48 Mb/s; ideally the algorithm should converge to 48 Mb/s.

*RRAA+*: To make RRAA converge, we propose a fix to the algorithm as shown in Algorithm 2 and call the new algorithm *RRAA+*. In brief, *RRAA+* maintains, for each data rate, the probability that it transitions to this rate from the next lower rate. Every time the loss at a rate exceeds the high threshold, the probability of returning to this rate is reduced before transitioning to the next lower rate. Also, every time the loss is lower than  $LT_k$ , the probability of the current data rate and the rates below is increased, assuming that all rates below the current rate can also be supported with the current channel conditions. Notice that the probability is only *used* on positive feedback (when the loss is less than  $LT_k$ ) for transitioning to a higher rate; transitions on negative feedback are deterministic. The multiplicative increase, multiplicative decrease (MIMD) parameters  $\alpha_1$  and  $\alpha_2$  are chosen such that the algorithm becomes stable. In our prototype,  $\alpha_1 = 2$  and  $\alpha_2 = 1.0905$ ; it takes eight increments to match one decrement.

<sup>3</sup>Frame is at the MAC level, whereas packet is at the driver level. Each packet can lead to multiple frame exchanges due to retransmissions.

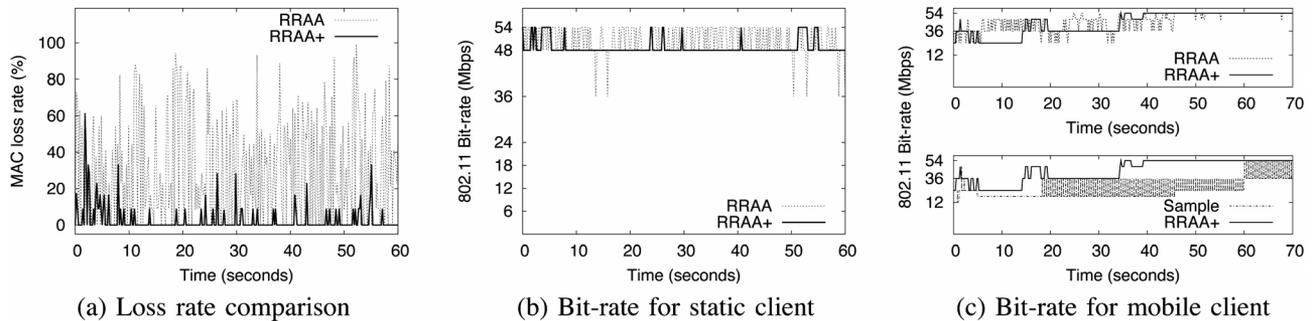


Fig. 7. Efficacy of RRAA+ over RRAA. (a) Loss rate comparison. (b) Bit-rate for static client. (c) Bit-rate for mobile client.

---

### Algorithm 2 RRAA+

---

```

1: if ( $loss > HT_k$ ) then
2:    $p[k]/ = \alpha_1$ 
3:    $k \leftarrow next\_lower\_rate$ 
4: else if ( $loss < LT_k$ ) then
5:   for (all rates  $j \leq k$ ) do
6:      $p[j]* = \alpha_2$ 
7:   end for
8:   if ( $rand() < p[next\_higher\_k]$ ) then
9:      $k \leftarrow next\_higher\_rate$ 
10:  end if
11: end if

```

---

To compare the performance of the algorithms, we implemented RRAA and RRAA+ in the Linux MadWifi driver. An implementation of SampleRate exists already in the MadWifi package. In our implementation, RRAA and RRAA+ is invoked every 200 ms or after 40 packets<sup>4</sup> have been received (defined as an *interval*), and the algorithms use loss rate estimated in the last interval for rate adaptation. We define the interval as above in order to be agile to user mobility and, at the same time, receive enough packets to reliably estimate channel conditions. In particular, VOIP calls generate 10 packets in the 200-ms interval, which is a reasonable number of samples. We now perform two experiments to demonstrate the benefit of RRAA+ over RRAA and SampleRate.

*Experiment 1:* We consider one sender and one receiver (two laptops with Atheros PCMCIA cards), with the receiver kept in NLOS with the sender at a distance of 15 m. We emulate a VOIP call between the nodes using the DITG traffic generator [10] that generates 50 packets per second with 20-byte payload, emulating a G.729.2 codec. We repeat the experiment with RRAA and RRAA+. Fig. 7(a) and (b) show the 802.11 frame loss and the data rate chosen by RRAA and RRAA+. The graphs show that RRAA incurs greater frame loss and, hence, increased number of retries because of not “learning” that 54 Mb/s rate is not appropriate for the link. In contrast, RRAA+ learns to avoid 54 Mb/s. Note that increased frame loss leads to reduced overall network throughput.

<sup>4</sup>Wong *et al.* [48] observe that 10–40 samples is a good enough number to reliably estimate the channel conditions.

---

### Algorithm 3 power\_control

---

```

1: if ( $(ERate_R - ERate_O > \tau_1)$  OR
    $(URTS_O > URTS_R)$  OR
    $(ETT_O - ETT_R > \tau_2)$ ) then
2:    $p[curpwr]/ = \beta_1$ 
3:    $curpwr \leftarrow next\_higher\_pwr$ 
4: else
5:   for (all pwr levels  $i \geq curpwr$ ) do
6:      $p[i]* = \beta_2$ 
7:   end for
8:   if ( $rand() < p[next\_lower\_pwr]$ ) then
9:      $curpwr \leftarrow next\_lower\_pwr$ 
10:  end if
11: end if

```

---

*Experiment 2:* We again consider one sender and one receiver, but we make the receiver mobile (by placing the laptop in a chair and dragging the chair). We set up the voice call as above again. The receiver starts at a distance of 40 m and in LOS with the sender and moves toward the sender to within a meter. The receiver is initially stationary and starts moving 15 s after the start of the experiment. We repeat this experiment for RRAA, RRAA+ and SampleRate. In Fig. 7(c), we observe that SampleRate takes longer to converge than both RRAA and RRAA+, while RRAA has the same oscillatory behavior as in Fig. 7(b). Such conservative rate selection by SampleRate leads to inefficient channel usage and reduces overall network throughput. Furthermore, as shown in [48], SampleRate can incorrectly lower the bit rate in response to collision induced losses.

In summary, we choose RRAA+ over others for rate adaptation in Symphony due to its three features—agility, convergence to the appropriate rate, and avoidance of rate adaptation because of collision-induced packet losses.

### C. Power Adaptation

Our goal for power adaptation is to tune each (sender, receiver) link in a WLAN to the lowest transmit power such that the performance metrics in the OPT phase are no worse than the corresponding metrics in the REF phase. Algorithm 3 shows the basic behavior of power adaptation in Symphony. The three conditions in line 1 detect undesirable rate adaptation, receiver-

side interference, and asymmetric channel access introduced by power control. Similar to RRAA+, the power control algorithm learns the lowest appropriate power level by maintaining the probability with which it should transition to a particular level. The algorithm executes once for every two intervals of the rate adaptation algorithm to adapt to user mobility. Furthermore, several rate adaptation intervals can occur in each of the REF and OPT phases, depending on traffic.

*Preventing Undesirable Rate Adaptation:* For detecting and preventing undesirable rate adaptation due to power control, for each link, the two contexts maintain an EWMA of the rate chosen by the rate control algorithm in response to the measured packet loss: for each  $rate_i$  chosen in interval  $i$ , we set  $ERate = ERate * \psi + rate_i * (1 - \psi)$  at the end of interval  $i$ . Everytime the power control algorithm is triggered, if the EWMA of rate in OPT phase ( $ERate_O$ ) is lower than that in the REF phase ( $ERate_R$ ) by a threshold  $\tau_1$ , transmit power is increased. In our implementation, the EWMA parameter  $\psi$  is chosen as 0.8. We choose  $\tau_1$  to be 3 Mb/s if  $ERate_R$  is above 48 or below 24 Mb/s, and 6 Mb/s otherwise. We make this choice because of the nonuniformity in 802.11 a/g bit rate granularity. Our idea is to place the threshold between the two consecutive rates.

*Preventing Receiver-Side Interference:* To detect that power control introduces receiver-side interference, we use the adaptive RTS/CTS mechanism. Similar to RRAA [48], RRAA+ includes a mechanism to detect if packet losses are happening due to collisions as opposed to degraded channel conditions. Our implementation of the adaptive RTS/CTS mechanism, however, differs significantly from RRAA. While RRAA was implemented on a per-frame basis (because of the availability of card firmware), we implement it on the basis of a window of packets, both for more reliable estimation of receiver-side interference and for obviating the need to modify firmware.

Using this mechanism, we maintain a performance metric—the utility of RTS (URTS)—that is set to 1 if the loss rate with RTS/CTS is less than the total loss rate in at least two out of the four last rate adaptation intervals, i.e., enabling RTS/CTS is helpful to reduce losses. Otherwise, URTS is set to zero. The rationale for waiting for four intervals is to determine the utility with greater reliability while trading off responsiveness. Furthermore, unless the receiver interference problem is sustained, we do not increase transmit power and let adaptive RTS address the problem. Now, if URTS is 0 in the REF phase and 1 in the OPT phase, it indicates that the receiver-side interference, which did not exist in the REF phase, was introduced by transmit power control. Line 1 in Algorithm 3 captures this condition and triggers a power increase.

*Preventing Asymmetric Channel Access:* To detect that power control introduces asymmetric channel access, we measure the EWMA of the ETT of each packet. The key idea here is that if a sender does not get a chance to transmit as frequently as in REF phase due to asymmetry in the OPT phase, the ETT in the OPT phase increases compared to the REF phase. If the ETT increases by more than a threshold  $\tau_2$ , we trigger power increase (as in Algorithm 3). In our implementation,  $\tau_2 = 100 \mu s$ .

If the interface card provides to the device driver the transmission time for each packet, EWMA of ETT can be easily cal-

culated. However, Atheros cards currently do not provide this information. Furthermore, multiple packets can be queued by the driver in the buffer of the card for efficiency, which makes ETT estimation nontrivial.

In our implementation, we overcome the above problem with the two functions in Algorithm 4 that exploit the interaction between the device driver and the interface card. We use the unique sequence numbers in packets that are sent to the card and keep one outstanding marked packet. We estimate when the packet transmission started using two variables:  $etime$ , which represents the time a marked packet was sent to the card, and  $itime$ , which represents the implicit time when the previous packet transmission was completed. We configure the card to provide per-packet transmission status feedback using `sysctl - wdev.wifi0.txintrperiod = 1`. This method ensures that ETT can be estimated even on packets that are buffered back to back and hence are not explicitly marked. Furthermore, we only consider packets that do not incur any retransmissions to reliably estimate the channel access delay, and consider packets sent at the same rate as in REF phase to avoid false positives due to rate-based ETT changes. Finally, since packet sizes affect the ETT because of the varied transmission time on the air, we maintain ETT in terms of 1500-byte packets. For smaller packets, the ETT is scaled to 1500 bytes before using it to calculate EWMA. To do so, we set for each packet  $fixedETT = ETT + (1500 - pkt\_size)/rate$ , where  $rate$  is the data rate used to transmit the packet. We validate experimentally that this approach of scaling ETT is reasonably accurate [39].

---

#### Algorithm 4 ETT Estimation

---

```

1: VARIABLES mark_seqno, itime, etime
2:
3: FUNCTION driver_send_pkt (seqno)
4:   if {(mark_seqno not set)}
5:     mark_seqno = seqno
6:     etime = curtime
7:   end if
8:
9: FUNCTION card_sent_pkt (seqno)
10:  if (mark_seqno = seqno) then
11:    ETT = curtime - MAX(itime, etime)
12:    itime = curtime
13:    unset mark_seqno
14:  else
15:    ETT = curtime - itime
16:    itime = curtime
17:  end if

```

---

*Granularity of Power Control:* Learning from our observations in Fig. 4, and given that power adaptation gets triggered at least twice in a second, in our implementation, Symphony increases and decreases power at a granularity of 3 dB, between [MIN\_PLEVEL, MAX\_PLEVEL]. This ensures agility to typical user mobility in WLANs. Furthermore, [47] observes that transmit power control at a finer granularity than 3 dB may not always be useful in indoor environments. The minimum and

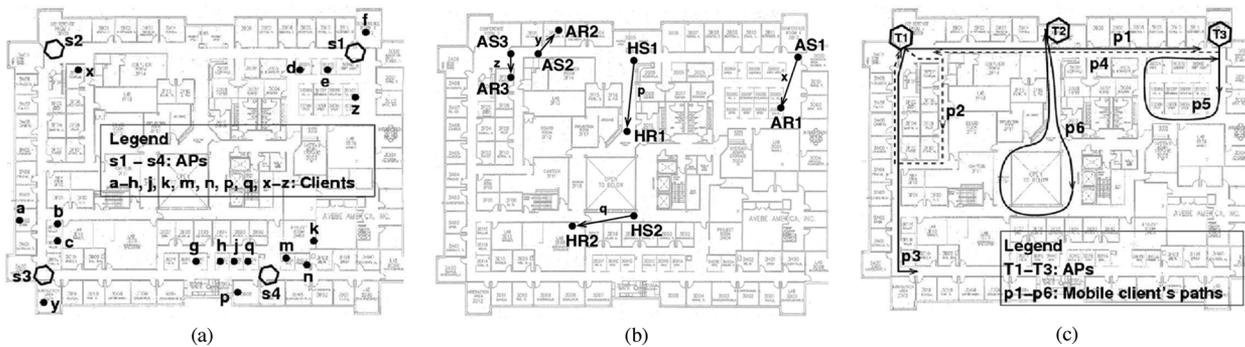


Fig. 8. Experimental setup. (a) Static nodes, (b) Problematic scenarios, (c) Mobile clients.

maximum transmit power values can be different on different 802.11 cards, vary with the frequencies used (such as in the 5-GHz band), and also vary based on the gains of the external antennas connected to the cards. However, we assume that the levels are discrete at a granularity of 3 dB. In our prototype with Atheros cards, we vary the power levels between 0 and 18 dBm.

The process of increasing and decreasing power is similar to rate adaptation in RRAA+. Symphony maintains for each power level, the probability that it transitions to this level from the next higher level. Every time at least one of the conditions on the performance metrics is satisfied, the probability of returning to this power level is reduced before transitioning to the next higher power. The MIMD parameters  $\beta_1$  and  $\beta_2$  are chosen to make the algorithm stable. In our prototype,  $\beta_1 = 3$  and  $\beta_2 = 1.14$ ; it takes eight increments to match one decrement. Again, the values for  $\beta_1$  and  $\beta_2$  are chosen to strike the tradeoff between the benefits of power control and stability of the algorithm; the above values are arrived at after experimenting with several scenarios.

With Symphony's approach of maintaining probability per power level, the transmit power of each sender will eventually converge to a point where the performance of each link is at least as good as in the REF phase. If the performance of a link at a given power level is similar to or better than that of the REF phase, the probability of returning to that power level and to any higher power level will converge to 1 due to the multiplicative increase of the probability. Otherwise, the probability of returning to that power level will converge to a small value due to the multiplicative decrease. To be responsive to changing channel conditions and mobility, we bound the probability to  $\frac{1}{64}$  in our implementation, such that it cannot be reduced below this value.

#### IV. EVALUATION

To demonstrate the achievement of the design goals outlined earlier, in this section, we carry out a systematic and extensive set of experiments in both controlled and uncontrolled environments. In what follows, we describe the experimental setups and then present our results.

##### A. Setup

Fig. 8 shows our setups for the different experiments. We broadly classify the setups into four categories: 1) *static nodes*

[Fig. 8(a)] in which clients are randomly placed in different cubicles and office rooms (represented by lowercase letters), and APs are placed close to the operational WLAN APs; 2) *problematic scenarios* [Fig. 8(b)] in which we carefully choose client and AP locations so as to emulate channel access asymmetry and receiver-side interference; 3) *mobile clients* [Fig. 8(c)] in which APs are placed close to operational WLAN APs (the AP is at location T1 for paths p1, p2, and p3; at T2 for path p6; and at T3 for paths p4 and p5), and clients associated with one of the APs move on different paths in the office building; and 4) the indoor ORBIT testbed [41]. For the setup with client mobility, the rationale behind the path selection is to approximate a reasonable mix of mobility patterns (moving toward/away from an AP), and channel conditions (LOS or NLOS) that occur in real situations.

In the first three setups, Symphony APs and clients are Dell laptops, while the ORBIT testbed uses custom-made, small form-factor PCs. All setups use Atheros cards (PCMCIA or mini-PCI), which transmit at a default (maximum) power level of 18 dBm<sup>5</sup> for 54 Mb/s data rate and allow Symphony to change transmit power at the desired granularity of 3 dB. The APs are connected by a 100-Mb/s wired network, and all nodes use the same 802.11 channel. While both clients and APs can execute Symphony, we perform most of our experiments with APs as senders and clients as receivers. This is because of the current limitation of Atheros cards that do not implement per-packet power control for ACK and CTS packets. To overcome this limitation for proper evaluation, unless specified otherwise, we run Symphony on APs and disable per-packet transmit power control on the clients. Instead, APs append transmit power information to each outgoing frame, and a slightly modified client driver extracts this information and sets the card to this power level (similar to how `iwconfig athN txpower$val sets power`). This ensures that all frames, including ACKs are returned at the configured power. Finally, except for the experiments that focus on interactions with Symphony-noncompliant nodes (e.g., legacy WLANs), all experiments are conducted on 802.11a channels to avoid interference from/to our office's WLAN that uses all three 802.11g channels (1,6,11).

For AP synchronization and configuration, instead of using a different central controller, we just use one AP as the master

<sup>5</sup>While the specifications mention a maximum transmit power of  $15 \pm 2$  dBm, we observe that the card can use up to 18 dBm.

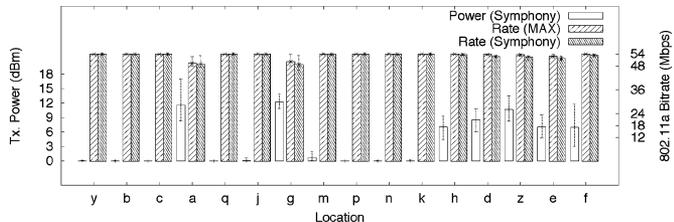


Fig. 9. Transmit power reduction for clients in several representative office locations.

AP and synchronize the others with the master through NTP on the wired network. Furthermore, for all the experiments in this section, we use 200 and 800 ms as the length of REF and OPT phase, respectively, and all APs are configured to start the REF phase at the beginning of each 1-s boundary. The choice of phase lengths is a tradeoff that ensures collecting reliable estimates in the REF phase while not compromising too much on energy savings [39].

*Two-Phase Synchronization:* To test that Symphony indeed executes synchronously on APs and clients, we performed micro-benchmarks with two APs and two clients—nodes make corresponding phase transitions within 5 ms of each other [39].

## B. Results

We now describe our results from several experiments on the above setups to demonstrate Symphony’s efficacy.

### Static Experiments

*Transmit Power Reduction:* To get a measure of the achievable transmit power reduction in near-typical indoor office environments, we consider all nodes in Fig. 8(a). We set up 10 1-min VOIP calls for each client. From the associated APs, the calls start at different times (separated by 5 min) for each client. The white bars in Fig. 9 show the average transmit power used by the Symphony APs for each client; in most typical user locations, the required transmit power can be substantially lower than the default 18 dBm. Furthermore, the error bars plot the minimum and maximum of the average transmit power per call, showing that the optimal power for maintaining a link’s performance can vary with time for even static locations. The other two bars show the rate chosen by Symphony and the rate when transmitting at maximum power. Symphony causes minimum effect on the data rate chosen. *In this setup, for the three cells with S1, S3, and S4 as APs, a per-cell solution would operate all links at the worst client’s transmit power, which is about 12 dBm. In contrast, Symphony enables 75%–80% of the clients in the cells to settle at 3–12 dB (i.e., 50%–94%) lower transmit power than 12 dBm.*

*Spatial Reuse:* Several studies demonstrate that transmit power control leads to increased spatial reuse [26], [35]. These works show simulation results over large topologies that cannot be easily realized in a prototype testbed. Here, we perform a small-scale spatial reuse experiment, where we consider links (S1, z), (S2, x), and (S3, y) as in Fig. 8(a) and make a subset of them operate simultaneously. We consider two cases: when the links operate at maximum power and when they operate

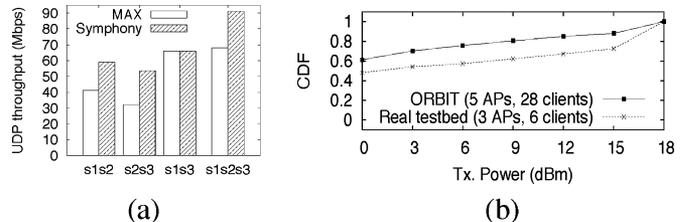


Fig. 10. Experiments demonstrating (a) increased spatial reuse and (b) significant power reduction even in a dense setup.

with Symphony. Fig. 10(a) shows the aggregate throughput of the links, and clearly demonstrates 30%–50% increased throughput for different combinations.

*Large-Scale Experiments:* To assess the effect of high node density on Symphony, we emulate dense deployments on the indoor ORBIT testbed [41] with five APs and 28 clients and on our office testbed with three APs and six clients [placed close to s4 in Fig. 8(a)]. The inter-AP distance is 5 m on ORBIT and 15 m in our office testbed. Clients are within 15 m of each AP in both cases. We setup bidirectional traffic between each client and its associated AP and enable Symphony on all nodes. Fig. 10(b) shows the CDF of the average transmit powers used by clients in each second, over a period of 60 s. We observe that Symphony enables clients to settle at much lower power levels in both sets of experiments. *For example, in the ORBIT experiment, clients settle at 0 dBm transmit power, and within 9 dBm, over 60% and 80% of the time, respectively.*

### Problematic Scenarios

*Avoiding Channel Access Asymmetry:* To demonstrate that Symphony avoids channel access asymmetry by intelligently adapting the transmit power of a link, we perform two sets of experiments. In the first set, we consider links x and y as in Fig. 8(b) (where AS1 and AS2 are senders and AR1 and AR2 are receivers) and set up backlogged UDP traffic to measure the link throughput. Link x always operates at the maximum power. We consider several cases—(i) links x and y running one at a time with y running Symphony; (ii) links x and y running simultaneously with y at a fixed 0 dBm; and (iii) links x and y running simultaneously with y running Symphony. Fig. 11(a) shows the results over 10 runs. In case (i), Symphony enables link y to operate at 0 dBm and still achieve full throughput. If link y is operated at 0 dBm together with link x, however, the throughput of link y drops significantly compared to x, as shown in case (ii) due to asymmetric channel access. We validate this by observing the difference in ETTs observed by links x and y. In case (iii), we show that Symphony increases the transmit power of link y to between 6 and 9 dBm to avoid asymmetry. Observe that link y did not have to operate at the maximum power to let link x perceive its transmission.

In the second set, we show that when there is inherent asymmetry between two links at default power, power control is beneficial in removing it and increasing the throughput and fairness of links. As in Fig. 8(b), a sender at location AS3 gets significantly affected when running in conjunction with a sender at AS1. For this experiment, we use links x and z and consider two cases: (i) both x and z running together at maximum power;

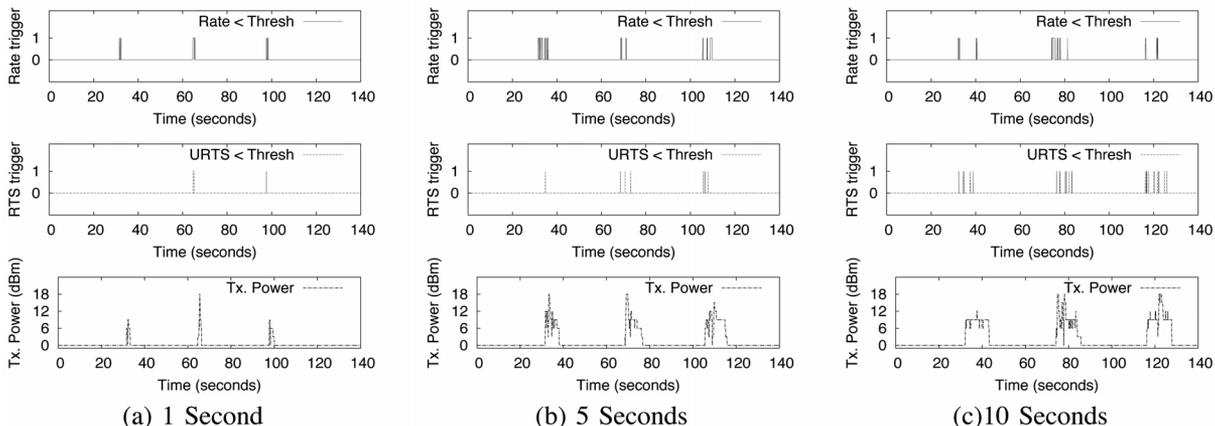


Fig. 12. Efficacy in avoiding receiver-side interference, which is introduced for (a) 1-, (b) 5-, and (c) 10-s durations at different times. Utility of RTS and/or a decrease in rate trigger increase in transmit power.

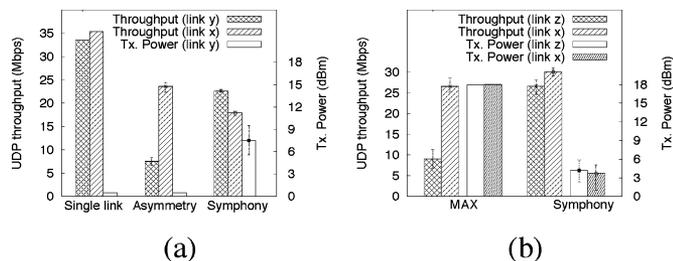


Fig. 11. Preventing asymmetric channel access. (a) Symphony’s ability to detect and avoid channel access asymmetry. (b) Power control removes inherent link asymmetry.

and (ii) x and z running together with Symphony. Results averaged over 10 runs in Fig. 11(b) show that, when operating at maximum power, link z gets significantly lower throughput than link x. When x and z use Symphony, both links achieve greater throughput because they are able to operate independently at lower transmit power, thereby also demonstrating increased spatial reuse due to power control. *In both sets of experiments, Symphony increases the throughput of asymmetry-affected links by 300%.*

**Avoiding Receiver-Side Interference:** To show that Symphony is effective in addressing receiver-side interference in mobile environments, we consider links p and q in Fig. 8(b). Link q operates at maximum power, whereas link p operates with Symphony. The setup is such that link p operates at 0 dBm when run individually, HS1 and HS2 share the channel at maximum power, and HS2 does not perceive transmission on link p if p operates at 0 dBm and hence destroys packets at HR1 (i.e., causes receiver-side interference). In each run, we start a 20-Mb/s UDP transfer on link p for 3 min, and start a 5-Mb/s transfer on link q for short periods (1, 5, and 10 s) of time, at different times during the 3 min.

The bottom graphs in Fig. 12 show that in response to link q’s entry and exit, Symphony on link p increases and decreases transmit power respectively to avoid the adverse effect of receiver-side interference. The graph shows that Symphony is responsive to receiver-side interference even at short timescales of 1 s. The top two graphs in each of the cases (a), (b), and (c) show

the rate and RTS triggers that identify the condition that triggered increased power as in line 1 of Algorithm 3. We note that the utility of RTS is not always sufficient to detect receiver-side interference, primarily because even RTS is sent at the chosen (lower) transmit power with Symphony, which may not be perceived by HS2. In such a case, rate drops and leads to increased transmit power, thereby letting RTS reach HS2. Recall that if RTS is useful, it helps avoid unnecessary reduction in rate.

*Mobility, Coexistence With Legacy, and Battery Life*

**Client Mobility:** To demonstrate Symphony’s agility to client mobility, we consider three AP locations and six client paths as shown in Fig. 8(c). On each of the paths, the client is mobile at a speed of 0.75 m/s. We again setup VOIP calls from the AP to the client. Fig. 13(a) shows for path p1 that moving the client away starts affecting the bit rate in the OPT phase, and hence Symphony increases transmit power to maintain the bit rate to the same level as in the REF phase. As the client moves farther, even the rate in REF phase falls. Fig. 13(b) shows that moving the client away on path p2 makes Symphony increase transmit power, but it also reduces the power when the client returns to the AP location. *Overall, Symphony opportunistically enables a link to operate up to 18 dB lower than the default (18 dBm).* Similar observations can be made on other paths [39].

For the same experiment, Fig. 13(c) shows the application-level loss rate for Symphony in comparison to using default maximum transmit power and the difference in R-score. R-score is a popular performance metric for the quality of voice calls [14]. An R-score of 70 or more is considered good voice quality. While R-score depends on several factors [14], in brief, the difference in R-score can be simplified to  $40 \times (\log(1 + 10e_m) - \log(1 + 10e_s))$ , where  $e_m$  and  $e_s$  represent the loss rate with maximum power and Symphony, respectively. The graph shows that Symphony incurs little extra impact on application-level loss on all the paths. Furthermore, *in the worse case (path 4), the R-score using Symphony deteriorates only by 3.4, and the average R-score deterioration using Symphony is 2.* While the actual R-score also depends on the end-to-end delay, we note that with even the maximum loss rate (as in path 4) using Symphony, in order to achieve

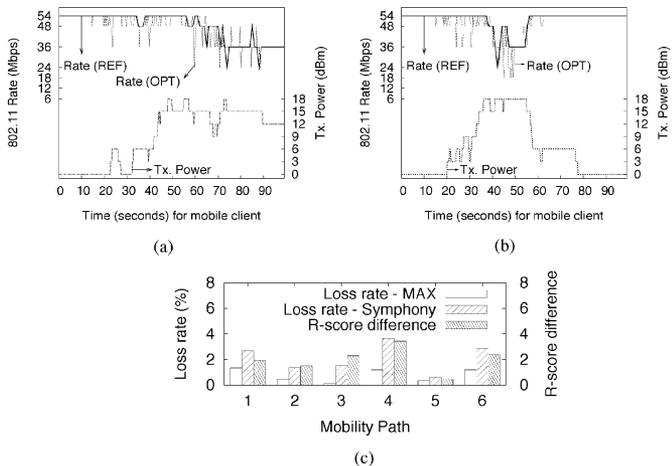


Fig. 13. Adaptation to mobility: **Symphony**'s behavior of rate and power in two different paths, and VoIP loss rate with mobility. (a) Path 1. (b) Path 2. (c) Application loss rate.

R-score as low as 70, the acceptable end-to-end delay is over 300 ms. Furthermore, in-depth analysis of the losses shows that significant part of the loss occurs when the client is far away from the AP. In a real mobility-enabled WLAN, mobile clients hand off to closer APs for better quality, and hence we believe that even this application-level loss will not occur in practice.

*Experiments With Noncompliant Nodes:* We now consider the case when **Symphony** is incrementally deployed in a setting where it has to operate in conjunction with noncompliant nodes such as other legacy WLANs. We first investigate whether transmit power reduction can be achieved in this situation. We consider the setup [Fig. 8(a)] with s3 as the AP location, and a,b,c,y as four client locations, with s3 running **Symphony**. We switch to channel 6 in 802.11g, on which there are other active APs and clients from our office WLAN. By observing beacons at the location s3, we determine that there are at least 13 non-compliant APs on channel 6.

For each client, the **Symphony** AP makes a 2-Mb/s transfer to every client every half-hour to estimate the power level appropriate for the clients. We perform this experiment over a period of 12 h mostly during regular office hours when the network is active. Fig. 14(a) shows the CDF of the average transmit power per client in each run, across several runs, for four clients. The graph shows that **Symphony** *opportunistically reduces the transmit power on all links even when operating in conjunction with a non-Symphony-compliant network*. For instance, links (s3, c) and (s3, y) operate 6 dB lower than the default transmit power 85% of the time.

We also study the effect of **Symphony** on noncompliant nodes in a controlled environment by conducting experiments on the ORBIT testbed. In this setup, we consider two APs and 11 clients using **Symphony** and two APs and 13 clients in legacy mode (using maximum transmit power and SampleRate for rate adaptation). For the legacy clients, we emulate realistic WLAN traffic (as described in [16]) while using two-way VOIP for the **Symphony** clients. We compare the average packet error rate (PER) for legacy clients with and without transmit power control on the **Symphony** part of the network. Fig. 14(b) shows the CDF of average PER per client. The graph shows that

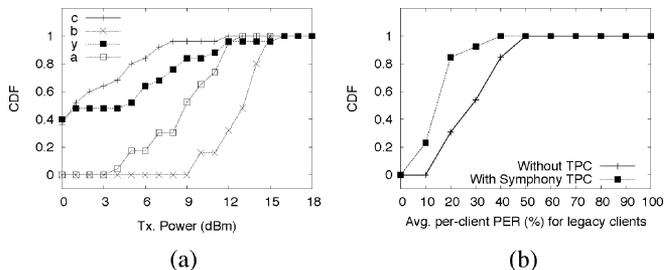


Fig. 14. Interaction with an operational network not running **Symphony**. (a) Transmit power reduction. (b) PER reduction.

TABLE III  
ENERGY SAVINGS WITH **Symphony** OVER MAX POWER

Experiment	Avg. Power (mW)	Energy saving over MAX (%)
Tx at 17dBm (MAX)	1548	
Tx at 0dBm (MIN)	1317	14.93
<b>Symphony</b> client : 2m from AP	1346	12.8
10m from AP	1382	10.6
20m from AP	1438	7.04
moving away from AP for 40s	1412	8.7

**Symphony** nodes have a positive effect of reducing PER for the legacy clients. Since the legacy clients use higher transmit power than the **Symphony** nodes, their transmissions more likely benefit from the *capture effect*, which in turn reduces the PER. Reduced PER results in the use of higher MAC bit-rates [39]. These results demonstrate **Symphony**'s incremental deployability.

*Potential Battery Life Improvements:* Due to lack of support for directly measuring the exact energy savings, we take a two-pronged approach to demonstrate the potential energy benefits of **Symphony**. First, we measure the energy consumption of a D-Link DWI-AG660 802.11 a/b/g PC card using a PC extender device, while executing **Symphony**, to show the achievable gains with existing cards. Second, we choose the analytical model in [37], and reference power consumption values from a newer power-optimized card (Broadcom's BCM 4328 card [1]), to demonstrate what may be achievable as technology improves in the future.

To measure the reduction in power consumption when using **Symphony** with existing PC cards, we plugged a D-Link PC card into a laptop via a PC extender device. This extender device exposes various pins to measure the current, and hence power, used by the card. We performed power measurements with **Symphony**, using a Tektronix TDS 3012 oscilloscope, for a set of static locations at different distances and with mobility. In our experiments, as traffic, we used backlogged UDP packets from the AP to the client. Our measurements in Table III show that a maximum of 12.8% energy can be saved relative to that used by MAX power. Note that, with the chosen card, the maximum possible energy savings is 14.93%.

With technology improvements, we foresee that higher gains may be achievable. To support this claim, we utilize the analytical model in [37]. Using the model, power consumption values from Broadcom's BCM 4328 card [1], as well as the CDF and the power dissipated by clients at each transmit power level, we can calculate the average power dissipated by **Symphony** clients

as  $P_d^S = 429$  mW for the ORBIT setup and  $P_d^S = 497$  mW for the indoor office setup. In contrast, transmitting at the maximum transmit power of 18 dBm would lead to a power dissipation of  $P_d^M = 766$  mW. We observe that the power consumption numbers at different transmit powers published for the Broadcom card [1] adhere well to the analytical model.

Since battery life is inversely proportional to the power dissipated, assuming that a card transmits and receives for equal amounts of time (and ignoring MAC protocol effects for simplicity and illustrative purposes), the active mode battery life improvement can be approximated as follows. At maximum transmit power, if  $\mathcal{L}_M$  is the battery life, and  $P_T$  is the power dissipated during reception, then

$$\mathcal{L}_M = \frac{K}{P_T + P_d^M + C} \quad (1)$$

where  $K$  is a proportionality constant and  $C$  represents the power consumed by the other components of the mobile (such as the processor, the graphics display, etc.), respectively. Similarly, if  $\mathcal{L}_S$  is the battery life with Symphony, then,  $\mathcal{L}_S = \frac{K}{P_T + P_d^S + C}$ . Then, the active-mode battery life improvement is  $\frac{\mathcal{L}_S - \mathcal{L}_M}{\mathcal{L}_M}$ . Substituting the  $P_d$  values calculated earlier and using the value for  $P_T = 295$  mW for the BCM 4328 card [1], we determine that the active-mode battery lifetime improvement with Symphony over using default transmit power can be 0%–46% for the ORBIT setup and 0%–33% for the office setup depending on the value of  $C$ . Compared to a per-cell power control solution [30], for Fig. 9, the improvement for each client can range from 0% to up to 26% with average being up to 17%. The higher the value of  $C$ , the lower will be the benefits of transmit power control. However, we envision that technology improvements will only reduce  $C$ 's value [42].

## V. DISCUSSION AND LIMITATIONS

While Symphony's per-link power control improves network throughput, its greedy approach may not achieve *maximum* network throughput. A network-wide optimization problem considering traffic profiles, locations of users, channel conditions for all users, etc., may lead to higher network throughput while compromising some users. However, such network-wide optimizations are also much harder to realize in practice, especially with user mobility.

Also, Symphony's current design is guided by the observation that reliable, per-link SINR feedback is not available in today's mobility-enabled WLANs. If such feedback were available in a timely manner, we expect that the rate adaptation component of Symphony will be simpler. However, the components for detecting hidden nodes and channel access asymmetry will remain.

For further discussion on the complementary nature of Symphony's TPC to clear channel assessment (CCA) tuning, on its need for synchronization, and its assumption on the prevalence of per-packet TPC in wireless cards, we refer the reader to our technical report [39].

## VI. CONCLUSION

In this paper, we propose Symphony, a novel synchronous two-phase rate and power control system for 802.11 WLANs. Symphony's agility in adaptation enables us to systematically address the key challenges facing per-link adaptive transmit power control while increasing network capacity and battery life of mobile devices. Through a detailed prototype study, we conclude that Symphony is: 1) effective in achieving the goals set forth; 2) easy to realize in a WLAN; 3) readily deployable even in the presence of noncompliant nodes while increasingly providing power control benefits as more nodes adhere to Symphony's strategy.

## REFERENCES

- [1] "Wireless LAN," Broadcom [Online]. Available: <http://www.broadcom.com/products/Wireless-LAN/802.11-Wireless-LAN-Solutions/BCM4328>
- [2] "Cisco Aironet 350 series access points," [Online]. Available: [http://www.cisco.com/univercd/cc/td/doc/product/wireless/airo\\_350/accesspts/ap350hig/ap350ch2.htm](http://www.cisco.com/univercd/cc/td/doc/product/wireless/airo_350/accesspts/ap350hig/ap350ch2.htm)
- [3] "Infonetics research," [Online]. Available: <http://www.infonetics.com/>
- [4] "Network time protocol," [Online]. Available: <http://www.ntp.org>
- [5] "Vo-Fi and dynamic power control: A possible mismatch?," [Online]. Available: <http://tinyurl.com/323at6>
- [6] "Where are the Wi-Fi phone power-save capabilities?," [Online]. Available: <http://tinyurl.com/3bcn4x>
- [7] "Atheros breaks the power barrier with new Wi-Fi solution for mobile devices," Oct. 2007 [Online]. Available: <http://atheros.com/news/AR6002.htm>
- [8] A. Akella, G. Judd, S. Seshan, and P. Steenkiste, "Self-management in chaotic wireless deployments," in *Proc. ACM MobiCom*, 2005, pp. 185–199.
- [9] W. Alliance, "WMM power save for mobile and portable Wi-Fi certified devices," Tech. Rep., 2005.
- [10] S. Avallone, S. Guadagno, D. Emma, A. Pescape, and G. Ventre, "D-ITG distributed internet traffic generator," in *Proc. QEST*, Sep. 2004, pp. 316–317.
- [11] J. C. Bicket, "Bit-rate selection in wireless networks," Master's thesis, MIT, Cambridge, MA, Feb. 2005.
- [12] P. Chevillat, J. Jelitto, and H. L. Truong, "Dynamic data rate and transmit power adjustment in IEEE 802.11 wireless LANs," *Int. J. Wireless Inf. Netw.*, pp. 123–145, 2005.
- [13] M. Chiang and J. Bell, "Balancing supply and demand of bandwidth in wireless cellular networks: Utility maximization over powers and rates," in *Proc. IEEE INFOCOM*, Mar. 2004, vol. 4, pp. 2800–2811.
- [14] R. Cle and J. Rosenbluth, "Voice over IP performance monitoring," *Comput. Commun. Rev.*, vol. 31, no. 2, pp. 9–24, Apr. 2001.
- [15] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *Proc. ACM MobiCom*, 2004, pp. 114–128.
- [16] M. A. Ergin, K. Ramachandran, and M. Gruteser, "Understanding the effect of access point density on wireless LAN performance," in *Proc. ACM MobiCom*, 2007, pp. 350–353.
- [17] G. J. Foschini and Z. Miljanic, "A simple distributed autonomous power control algorithm and its convergence," *IEEE Trans. Veh. Technol.*, vol. 42, no. 4, pp. 641–646, Nov. 1993.
- [18] P. Hande, S. Rangan, and M. Chiang, "Distributed uplink power control for optimal sir assignment in cellular data networks," in *Proc. IEEE INFOCOM*, 2006.
- [19] I. Haratcherev, K. Langendoen, R. Legendijk, and H. Sips, "Hybrid rate control for IEEE 802.11," in *Proc. Int. Workshop Mobility Manag. Wireless Access*, 2004, pp. 10–18.
- [20] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *Proc. ACM MobiCom*, Rome, Italy, 2001, pp. 236–251.
- [21] J. Jeong, D. E. Culler, and J.-H. Oh, "Empirical analysis of transmission power control algorithms for wireless sensor networks," in *Proc. IEEE INSS*, 2007, pp. 27–34.
- [22] E.-S. Jung and N. H. Vaidya, "A power control MAC protocol for ad hoc networks," in *Proc. ACM MobiCom*, 2002, pp. 36–47.
- [23] A. Kamerman and L. Monteban, "WaveLAN-II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Tech. J.*, pp. 118–133, 1997.

- [24] V. Kawadia and P. R. Kumar, "Principles and protocols for power control in ad hoc networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, pp. 76–88, 2005.
- [25] J. Kim, S. Kim, S. Choi, and D. Qiao, "CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs," in *Proc. IEEE INFOCOM*, 2006.
- [26] T.-S. Kim, H. Lim, and J. C. Hou, "Improving spatial reuse through tuning transmit power, carrier sense threshold, and data rate in multihop wireless networks," in *Proc. ACM MobiCom*, 2006, pp. 366–377.
- [27] M. Lacage, M. H. Manshaei, and T. Turletti, "IEEE 802.11 rate adaptation: A practical approach," in *Proc. MSWiM*, 2004, pp. 126–134.
- [28] S. Lin, J. Zhang, G. Zhou, L. Gu, J. A. Stankovic, and T. He, "ATPC: Adaptive transmission power control for wireless sensor networks," in *Proc. SenSys*, 2006, pp. 223–236.
- [29] "Multiband Atheros driver for WiFi," MadWifi [Online]. Available: <http://www.madwifi.org/>
- [30] V. Mhatre, K. Papagiannaki, and F. Baccelli, "Interference mitigation through power control in high density 802.11 WLANs," in *Proc. IEEE INFOCOM*, 2007, pp. 535–543.
- [31] J. Monks, V. Bharghavan, and W.-M. Hwu, "A power controlled multiple access protocol for wireless packet networks," in *Proc. IEEE INFOCOM*, 2001, vol. 1, pp. 219–228.
- [32] A. Muqattash and M. Krunz, "Power controlled dual channel (PCDC) medium access protocol for wireless ad hoc networks," in *Proc. IEEE INFOCOM*, 2003, pp. 470–480.
- [33] A. Muqattash and M. Krunz, "POWMAC: A single-channel power-control protocol for throughput enhancement in wireless ad hoc networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 5, pp. 1067–1084, 2005.
- [34] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar, "Power control in ad hoc networks: Theory, architecture, algorithm and implementation of the COMPOW protocol," in *Proc. Eur. Wireless Conf.*, 2002, pp. 156–162.
- [35] V. Navda, R. Kokku, S. Ganguly, and S. Das, "Slotted symmetric power control in managed WLANs," NEC Laboratories America, Tech. Rep. [Online]. Available: <http://www.nec-labs.com/~ravik/RESEARCH/contour.pdf>
- [36] J. Pavan and S. Choi, "Link adaptation strategy for IEEE 802.11 WLAN via received signal strength measurement," in *Proc. IEEE ICC*, May 2003, vol. 2, pp. 1108–1113.
- [37] D. Qiao, S. Choi, A. Jain, and K. G. Shin, "Miser: An optimal low-energy transmission strategy for IEEE 802.11a/h," in *Proc. ACM MobiCom*, 2003, pp. 161–175.
- [38] D. Qiao, S. C. A. Jain, and K. G. Shin, "Adaptive transmit power control in IEEE 802.11a wireless LANs," in *Proc. IEEE VTC*, Apr. 2003, vol. 1, pp. 433–437.
- [39] K. Ramachandran, R. Kokku, H. Zhang, and M. Gruteser, "Synchronous two-phase rate and power control in 802.11 WLANs," [Online]. Available: <http://www.nec-labs.com/~ravik/RESEARCH/symphony-tech.pdf>
- [40] K. Ramachandran, H. Kremo, M. Gruteser, P. Spasojevic, and I. Seskar, "Scalability analysis of rate adaptation techniques in congested IEEE 802.11 networks: An orbit testbed comparative study," in *Proc. WoWMoM*, 2007, pp. 1–12.
- [41] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh, "Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols," in *Proc. IEEE WCNC*, Mar. 2005, vol. 3, pp. 1664–1669.
- [42] C. W. S. Gaudin, "Intel unveils new low-power chip/building block," [Online]. Available: <http://tinyurl.com/4wxdr>
- [43] V. Shah, E. Gelal, and S. V. Krishnamurthy, "Handling asymmetry in power heterogeneous ad hoc networks," *Comput. Netw.*, vol. 51, no. 10, pp. 2594–2615, 2007.
- [44] A. Sheth and R. Han, "A mobility-aware adaptive power control algorithm for wireless LANs," presented at the IEEE CAS Low Power Workshop, 2002.
- [45] A. Sheth and R. Han, "An implementation of transmit power control in 802.11b wireless networks," Dept. Comput. Sci., Univ. Colorado, Tech. Rep., 2002.
- [46] A. Sheth and R. Han, "SHUSH: Reactive transmit power control for wireless MAC protocols," in *Proc. WICON*, 2005, pp. 18–25.
- [47] V. Shrivastava, D. Agarwal, A. Mishra, S. Banerjee, and T. Nadeem, "Understanding the limitations of transmit power control for indoor WLANs," in *Proc. IMC*, 2007, pp. 351–364.
- [48] S. H. Y. Wong, S. Lu, H. Yang, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *Proc. ACM MobiCom*, 2006, pp. 146–157.
- [49] Y. Agarwal, R. Chandra, A. Wolman, P. Bahl, K. Chin, and R. Gupta, "Wireless wakeups revisited: Energy management for VOIP over Wi-Fi smartphones," in *Proc. ACM MobiSys*, 2007, pp. 179–191.
- [50] J. Zhu, X. Guo, S. Roy, and K. Papagiannaki, "CSMA self-adaptation based on interference differentiation," in *Proc. IEEE GlobeCom*, 2007, pp. 4946–4951.



**Kishore Ramachandran** received the Bachelor's degree in electronics and telecommunications engineering from the University of Mumbai, Mumbai, India, in 2000; the Master's degree in computer and systems engineering from Rensselaer Polytechnic Institute, Troy, NY, in 2002; and the Ph.D. degree in computer engineering from WINLAB at Rutgers University, New Brunswick, NJ, in 2009.

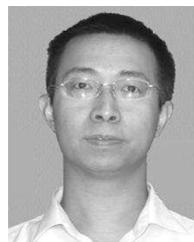
He is a Research Staff Member with NEC Laboratories America, Inc., Princeton, NJ. His research interests broadly include measurement, design, and analysis of wireless systems and networks. His work spans Internet congestion control, wireless network testbed design, resource management for WiFi networks, and vehicular networks. Past projects include building the world's largest, open-access, indoor wireless testbed (ORBIT), which was awarded the NSF's 4th Annual Schwarzkopf Award for Technological Innovation in 2008.



**Ravi Kokku** received the B.Tech. degree in computer sciences from the Indian Institute of Technology, Kharagpur, India in 1997 and the M.S. and Ph.D. degrees in computer sciences from the University of Texas at Austin in 2004 and 2005, respectively.

He is a Research Staff Member with NEC Laboratories America, Princeton, NJ. From 1997 to 2000, he worked at Hughes Software Systems, Gurgaon, India, on the design and development of satellite communication systems. His research interests include designing and prototyping innovative systems in domains including WiFi and WiMAX networks, wide-area data replication, and multiservice routers.

Dr. Kokku is a recipient of several awards, including the IBM doctoral fellowship (2001–2004) and Best Paper Award at BROADNETS 2008.



**Honghai Zhang** received the B.S. degree from the University of Science and Technology of China, Hefei, China, in 1998, and the Ph.D. degree from the Department of Computer Science, University of Illinois at Urbana-Champaign, in 2005.

He is currently a Research Staff Member with NEC Laboratories America, Princeton, NJ. Prior to his current position, he was in the Wireless Advanced Technology Laboratory of Alcatel-Lucent, Whippany, NJ. His research interests include scheduling, interference mitigation, and video streaming in mobile wireless networks and cognitive radio networks.

Dr. Zhang was a recipient of Vodafone Fellowship during his Ph.D. study.



**Marco Gruteser** received the Vordiplom from Darmstadt University of Technology, Darmstadt, Germany, in 1998 and the M.S. and Ph.D. degrees from the University of Colorado at Boulder in 2000 and 2004, respectively, all in computer science.

He is an Assistant Professor of electrical and computer engineering with WINLAB, Rutgers University, New Brunswick, NJ. He was also a Research Associate at the IBM T. J. Watson Research Center, Hawthorne, NY, from 2000 to 2001. His research interests include location-aware networking, the design of location privacy techniques, and applications in vehicular networks.

Prof. Gruteser is a recipient of the National Science Foundation CAREER Award. He also received the Schwartzkopf Prize for Technological Innovation as a member of the ORBIT wireless testbed team. He has served on the program committee of numerous conferences, including the ACM MobiSys and the IEEE INFOCOM.