

Implications of Power Control in Wireless Networks: A Quantitative Study

Ioannis Broustis*, Jakob Eriksson, Srikanth V. Krishnamurthy, and Michalis Faloutsos

Department of Computer Science and Engineering,
University of California, Riverside
{broustis, jeriksson, krish, michalis}@cs.ucr.edu

Abstract. The use of power control in wireless networks can lead to two conflicting effects. An increase in the transmission power on a link may (i) improve the quality and thus the throughput on that link but, (ii) increase the levels of interference on other links. A decrease in the transmission power can have the opposite effects. Our primary goal in this work is to understand the implications of power control on interference and contention. We conduct experiments on an indoor mesh network. Based on analysis of our experimental data, we identify three interference scenarios: a) the overlapping case, where the aggregate throughput achievable with two overlapping links cannot be improved via power control; b) the hidden terminal case, where proper power control can primarily improve fairness and, c) the potentially disjoint case, where proper power control can enable simultaneous transmissions and thus improve throughput dramatically. We find that power control can significantly improve overall throughput as well as fairness. However, to our surprise, we note that using virtual carrier sensing in conjunction with power control generally degrades performance, often to a large degree.

Keywords: Wireless Networks, Interference, Carrier Sensing, Experimentation, Network Topology, Testbed.

1 Introduction

The goal of this paper is to characterize interference effects in IEEE 802.11-based wireless mesh networks and examine the impact of power control on interference. An increase in transmission power may result in: (a) increased quality of reception and hence, potentially higher throughputs at the intended receiver, and (b) increased interference levels. These two effects are conflicting in terms of providing the best network-wide throughput. We experimentally evaluate power control as a means of improving wireless network performance. We also evaluate the use of virtual carrier sensing (RTS/CTS) in conjunction with power control. In this study, we focus on the interference between pairs of links, and provide detailed experimental results for a wide variety of such pairs.

Based on our experimental results, we identify three interference scenarios. a) The *overlapping* case: if two links are overlapping, neither the use of power control nor the use of RTS/CTS messages can help improve the aggregate throughput. If two links

* This work is supported in part by the NSF CAREER Grant No. 0237920 and the NSF NRT grant No. 0335302.

overlap, only one of the links can be active at any given time. Thus, the maximum total achievable throughput would be the maximum throughput achievable on one of these links. b) The *hidden-terminal* case: in this case, we find that proper power control is essential for ensuring fairness. To our surprise, we observe that the use of RTS/CTS in conjunction with power control generally results in degraded overall throughput and fairness in the hidden-terminal case. c) The *potentially disjoint* case: here, we find that power control can result in dramatic performance improvements. However, the use of virtual carrier sensing tends to remove opportunities for simultaneous transmissions, and consequently results in significantly lower overall throughput. We discuss these three cases in detail and explain the observed behavior.

Our work in perspective. Our results are based on measurements on an indoor testbed. Outdoor environments are subject to different constraints and may yield different results. We also limit our experiments to pairs of links. While we expect our results to carry over to multi-link interference, further studies are needed to verify this. While we do not propose an online power control algorithm, our results clearly demonstrate the need for an adaptive power control mechanism with the IEEE 802.11 MAC protocol. We believe that our study will stimulate further research on power control in a variety of settings, as well as provide insight towards the design of an adaptive power control mechanism. We wish to point out that while there are many studies on power control, they have been based predominantly on simulation or analysis. To the best of our knowledge, there have been no extensive experimentation studies on power control and interference in conjunction.

In section 2 we provide background on carrier sensing and discuss related studies on power control and interference. Section 3 describes our testbed and some initial experiments used for validation purposes. In section 4 we describe our experimental methods and present results related to interference effects. Section 5 concludes the paper.

2 Background and Previous Work

We revisit the IEEE 802.11 MAC protocol and in particular physical and virtual carrier sensing. We also discuss related studies and how our work differs from these efforts.

The Distributed Coordination Function (DCF) in the IEEE 802.11 MAC uses the CSMA/CA algorithm. In particular, this algorithm hinges on both physical and virtual carrier sensing. With physical carrier sensing, in order to avoid collisions, a node that intends to transmit first senses the medium. If an ongoing transmission is detected, the node tries again after a back-off period that is specified by the IEEE standard [2]. If no traffic is detected, the node proceeds with its transmission. However, physical carrier sensing suffers from the presence of *hidden* and *exposed* terminals. First, two transmitters may be unable to sense each others' carriers and hence, their transmissions may collide at the intended receiver. In this case, the two transmitters are said to be *hidden* from each other. Second, a transmitter (say A) might suppress its transmission since it detects a carrier from another transmitter (say B); however, this may be overly conservative, as simultaneous transmissions are sometimes feasible despite A hearing B's carrier. In this case, transmitter A is said to be *exposed* to B's transmission. Virtual carrier sensing in IEEE 802.11, using RTS (Request-To-Send) and CTS (Clear-To-Send)

control messages, is intended to address the hidden terminal problem, and has been known to exacerbate the exposed terminal problem. More details on the DCF function may be found in [2].

Next, we briefly discuss relevant previous work on interference characterization and power control in 802.11-based networks. As with our work, these efforts assume omnidirectional communications. Power control in mesh networks has received a lot of attention since it is attractive in terms of providing energy savings and spatial reuse. However, most efforts assume that the network can be modeled using unit disk graphs (UDGs) [16,6,17,11]. There have been more recent efforts that account for more realistic channel models [10,8]. However, most of these efforts rely on simulations and analyses; the models used may not be appropriate in all settings although they do provide better representations of the wireless channel. In [10], Muqattash and Krunz, design a power controlled MAC protocol that allows nodes to adjust their transmission powers while allowing for some interference margin at the receiver. In [7], distributed power control algorithms are designed; the algorithms take into account node sensitivities to current interference levels.

Experimental studies on wireless networks have recently gained popularity. Studies of wireless mesh networks in [5] and [4] demonstrate that the popular unit disk graph models are unlikely to hold in real networks. However, these efforts do not characterize interference between links nor do they consider power control.

Akella et al. [1] use large-scale measurement data from various cities to show how common it is to have tens of APs deployed in close proximity of each other. The paper proposes a power control and rate adaptation scheme (PERF) for the purpose of reducing interference among neighbors. We discuss PERF further in section 4. Padhye et al. [13] also study the problem of estimating pairwise interference among links in a multi-hop wireless testbed. They propose a link interference estimation method and study it in a variety of settings. Sinha et al., in [18], perform experiments to observe the differences between unicast and broadcast link properties. To study the impact of interference, they consider various distance-related interference scenarios on a grid indoor-testbed. Son et al. [15] study the effects of transmission power control on individual wireless link quality. They focus on sensor networks and perform their experiments on low-power RF links (up to 10 dBm). The interference effects are not studied.

3 Experimental Setup

In this section we describe our experimental testbed, and present initial experimental results on the stability and power controlled behavior for isolated links. These results provide a basis for our studies on the interference effects.

Our indoor testbed is comprised of 15 Soekris net4826 nodes [14], deployed in the 3rd floor of Engineering Building Unit II at the University of California, Riverside; the network is depicted in figure 1. Each node runs a Debian v3.1 Linux distribution with kernel version 2.6.13.2 and mounts its root partition over NFS from a server at start-up. We have equipped nodes with EMP-8602-6G 802.11a/b/g WiFi cards, which embed the Atheros AR5006 chipset [3]; the cards are controlled by the Linux MadWifi driver (version 0.90) [9]. Each card is connected to a 5-dBi gain, external omnidirectional antenna.



Fig. 1. Our indoor-testbed deployment. Nodes are represented by dots along with their IDs.

We use the 802.11a mode in order to avoid interference from co-located 802.11b/g networks; our testbed is the only 802.11a network in the area. For our experiments, we use channel 56, which corresponds to 5.28 GHz. The transmission power is varied between 1 and 16 dBm in our experiments. We use SampleRate [4,5] as the rate control algorithm. In order to derive reliable measurements, we run multiple experiments between many different pairs of links, and for different time periods of the day. We use 30 seconds of back-to-back 1500-byte UDP packets as our traffic load.

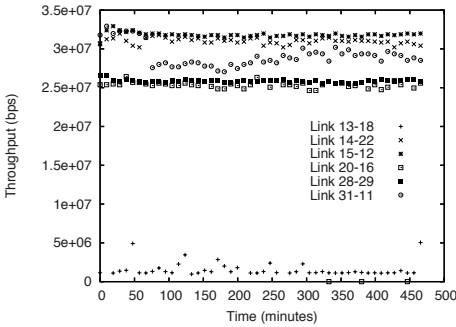


Fig. 2. Links exhibit relatively small variability over time

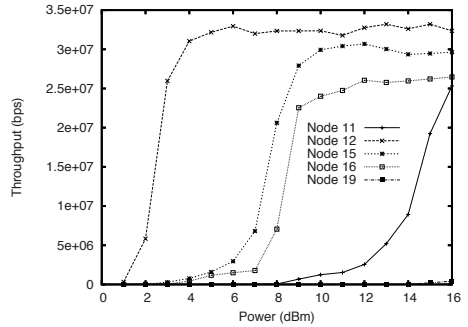


Fig. 3. Throughput (bps) vs. power (dBm) for the links of node 20

Reliability of Results. In order to determine the stability of our results, we ran a large number of 30-sec experiments, considering varying powers on a large set of links, each of which was activated in isolation. Figure 2 depicts the throughputs observed for 6 representative links, for 16 dBm of power. We observe that there is only a minor variation in performance, for most of the links under study. The throughput on the link 13→18 exhibits higher variance than the other links. This is because this link is relatively poor, as one can see from the achieved throughput. This set of experiments demonstrates the stability of our experimental results over time.

Throughput vs. Transmission Power. We observe the quality of isolated individual links (only one link is active at any given moment), for varying transmission power

settings. As shown in figure 3, the throughput achieved on any given link depends on the transmission power used. We also observe that the maximum throughput for a link is achieved at a certain power threshold. For example, for the link 20→12 we observe that maximum throughput is achieved with about 4 dBm of transmission power. The throughput saturates here and a further increase in the transmission power does not significantly increase performance. Thus, isolated links do not always have to transmit with maximum power, in order to achieve maximum throughput. One may attribute this to the limited modulation and encoding possibilities with IEEE 802.11a.

4 Effect of Transmission Power on Interference and Contention

The overarching objective of our experimental study is to understand the transmission power trade-off: increased reliability and performance on one link vs. increased levels of interference caused to other links.




Fig. 4. Using the maximum transmission power does not result in the optimal throughput

In our experiments, we activate two links at a time. We then study the throughput achieved by each link as we vary their respective transmission powers. The transmitters of these links send bulk UDP traffic to their respective receivers. The transmission power is selected between 1 and 16 dBm, and ACKs are transmitted at the same power as DATA packets. Experiments have a 30-second duration, during which transmission power is maintained at a fixed level. After the expiration of this time period, each link is configured with a new power setting, and the experiment is repeated. We use the same carrier sensing threshold for all experiments (the default setting in our cards). We utilize the *iperf* measurement tool (version 2.0.2) [12] to measure link performance. We consider cases where we (i) disable or (ii) enable virtual carrier sensing with RTS/CTS.

4.1 Achievable Performance Gains

We would like to determine by how much, if at all, performance can be improved through the use of an optimal power control mechanism. Since such a power control mechanism is not available, we resort to exhaustive search. Figure 4 depicts the maximum total throughput observed for a set of link pairs when using the optimal power settings, as well as when using the maximum transmission power. Both the cases of having




Fig. 5. A pair of overlapping links. Power control cannot improve the performance.

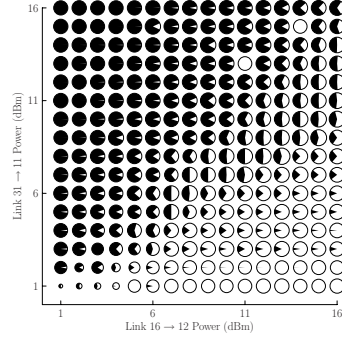


Fig. 6. A pair of links with a hidden terminal problem. Power control necessary for fairness.




Fig. 7. A pair of potentially disjoint links. Power control can improve throughput up to 100%.

RTS/CTS enabled or disabled are plotted. We observe that using maximum power does not necessarily result in maximum network throughput. Thus, an opportunity exists for significantly improving network performance through power control. Moreover, using RTS/CTS does not appear to improve performance for any of the link pairs in question. We address this in more detail below. Next, we take a closer look at the collected data, to see the effect of transmission power settings on throughput as well as fairness.

4.2 Three Types of Interference Behavior

The above experiments were run on a set of 16 representative link pairs. Analyzing the experimental data in more detail, we observed that each of the 16 pairs could be categorized as belonging to one of three general cases. The three interference cases identified were **overlapping** (5 links out of 16), **hidden-terminal** (9 links out of 16) and **potentially disjoint** (2 links out of 16).




Fig. 8. Achieved throughputs for the overlapping case, when RTS/CTS is enabled




Fig. 9. RTS/CTS worsens performance and fairness in the hidden terminal case




Fig. 10. Enabling RTS/CTS degrades performance in the disjoint case

We illustrate the three different types using the **throughput matrices** of representative link pairs (figures 5, 6 and 7). A throughput matrix illustrates the throughput achieved by each link, for all combinations of transmission powers. In figure 5, the horizontal and vertical axes represent the powers used by each link, respectively. The area of a disc corresponds to the sum total throughput achieved. The black and white portions of each disc correspond to the throughputs of the links indicated on the vertical and horizontal axes, respectively. For example, in figure 6, at coordinate (11,8), good throughput is achieved, and each link received exactly half of that throughput (17.5 Mbits/sec each). At coordinate (11,16), total throughput was again good, but link 31 → 11 received most of that throughput (30.4 Mbits/sec), leaving link 16 → 12 with very little (2.55 Mbits/sec). We will now describe each case in more detail, and study its effects in the corresponding throughput matrix.

Overlapping. In this case, the two links always contend, irrespective of the power settings used. Figure 5 shows an example of such a case. This case typically appears where the channel between the two senders is better than the channel from at least one sender

to its corresponding receiver. In this example, nodes 12 and 20 are the transmitters and nodes 11 and 15 are the receivers, respectively (figure 1). The link 12→11 is a poor link and node 12 needs to transmit with at least 12 to 13 dBm to achieve a reasonable throughput. Similarly, node 20’s power has to be at least 8 to 9 dBm. Transmitters 20 and 12 are within line-of-sight of each other (as seen in figure 1) and the link between them is very good. As a result, they can hear each others’ transmissions and back-off accordingly. Hence, only one of the two links will carry traffic at a time, and the maximum total throughput of the two links will never be higher than the maximum of what is achievable on the best of the two links.

Hidden Terminal. The majority of the considered link pairs belong to this type. Here, senders cannot detect each others’ transmissions and physical carrier sensing fails. Multiple transmissions arrive simultaneously at a receiver. This is not necessarily a problem if the signal strength of the desired transmission is significantly higher than that of the interfering transmission. This competition between signal and interference is observed in figure 6. As node 16 increases its power, it causes higher levels of interference on the link 31→11; collisions occur at node 11. However, as node 31 increases its transmission power, its signal strength increases. Hence node 11 is now able to decipher node 31’s data packet. Note that while total throughput varies little with transmission power, if this power is over a certain level, fairness is consistently better along a diagonal in the throughput matrix. Power settings along this diagonal optimally balance signal quality and interference between the two links.




Fig. 11. Throughput matrix for a second link pair suffering from a hidden terminal problem, without RTS/CTS




Fig. 12. Throughput matrix for the second pair, with RTS/CTS. Though RTS/CTS in some cases improves fairness, this is at the cost of severely degraded overall throughput.

Potentially Disjoint. Potentially disjoint links are links on which, accurate power settings enable simultaneous successful transmissions, resulting in up to a doubling of the overall throughput. Potentially disjoint links are characterized by peaks in the throughput matrix exceeding the maximum capacity of a single link. This is depicted in figure 7, for the links 15→16 and 22→31. These two links are topologically well-separated

(figure 1) and are of high quality (there is a line-of-sight path between the transmitters and their corresponding receivers). The maximum total throughput (approximately twice the capacity of a single link) is achieved when link 15→16 uses power 9 dBm and when link 22→31 uses 6 dBm. This suggests that when the transmission powers are set to the above values, the two links can operate interference-free. As links increase their powers, we observe that the total throughput decreases.

From these observations it is apparent that in order to achieve high total network throughput, appropriate setting of the transmission power is essential. In other words, an online power adaptation mechanism is needed that can, for each link, determine the power that is the most appropriate for use at a given instance in time. One proposed scheme, PERF [1], partially addresses this issue. PERF gradually decreases transmission power as long as throughput is not affected. PERF can be expected to successfully address the potentially disjoint case, correctly reducing power until spatial reuse is made possible. However in the hidden terminal case, PERF will use the maximum power, resulting in suboptimal fairness conditions. While PERF is a step in the right direction, we conclude that further improvements are necessary.

4.3 Use of RTS/CTS in Indoor Environments

Finally, we investigate how adding the RTS/CTS handshake affects interference and contention characteristics. We repeat the previous experiments, but with RTS/CTS enabled. Given that the purpose of RTS/CTS is to address the hidden terminal problem, we expect the results to somewhat improve in the hidden terminal cases. From the measurements we observe the following:

Overlapping. The performance of overlapping links is slightly worsened due to RTS/CTS overhead (figure 8).

Hidden Terminal. We were surprised to note that RTS/CTS consistently underperformed plain vanilla CSMA in *all* of the hidden terminal cases we tested. Figure 9 shows the throughput matrix for a typical hidden terminal case. Overall throughput is consistently lower when RTS/CTS is enabled. Fairness was affected to an even larger degree, to the extent that the stability of our results was significantly impacted. Note that there exist cases where, for some power settings, the use of RTS/CTS resulted in improved fairness. However, this improvement is generally offset by a large reduction in overall throughput (see figures 11, 12). We suspect that this is due to RTS/CTS exacerbating the exposed terminal problem, reducing the opportunity for parallel transmissions.

Potentially Disjoint. Finally, links in a potentially disjoint case are also negatively impacted by RTS/CTS (see figure 10). While overall throughput can still reach levels higher than the capacity of a single link, it is significantly lower than what was achieved with plain vanilla CSMA. The *regime* of powers in which spatial reuse is achieved is also reduced in size, requiring more precise power control, and generally higher transmission powers (compare to figure 7).

In conclusion, RTS/CTS appears to be entirely detrimental to the performance of indoor wireless networks, if power control is available. However, additional experiments need to be performed, in a wider range of environments and with a variety of hardware, before any firm conclusions be drawn.

5 Conclusion

While power control can affect the quality of transmission on a given link, the used power dictates the interference projected on other links and thus, can also affect the performance on other links. We perform measurements to quantify this trade-off with power control in an indoor experimental network with IEEE 802.11a nodes, both with and without virtual carrier sensing. We identify three types of interference behavior. In the *overlapping* case, power control and/or RTS/CTS do not increase the maximum achievable throughput. In the *hidden terminal* case, power control alleviates the hidden terminal problem and thus, improves the throughput. However, the use of RTS/CTS in conjunction with power control consistently degrades both the overall throughput and fairness. In the *potentially disjoint* case, power control can help in activating the links simultaneously and thereby can yield almost a twofold increase in the achievable throughput as compared to a case where the default maximum power is used. Again, virtual carrier sensing exacerbates the exposed terminal problem, significantly degrading overall performance.

We used IEEE 802.11a nodes, to avoid interference from co-located 802.11b/g networks. Even though results with 802.11b/g may differ for the links under investigation, our conclusions regarding the observed types of interference will still hold. Overall, our studies suggest that accurate power control holds great promise for improving the performance of indoor wireless networks.

References

1. A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self-management in chaotic wireless deployments. In *MOBICOM*, 2005.
2. ANSI/IEEE802.11-Standard. 1999 edition.
3. Atheros/AR5006chipset. <http://www.atheros.com/pt/ar5006bulletins.htm>.
4. J. Bicket. Bit-rate selection in wireless networks. In *Master's Thesis, Department of Electrical Engineering and Computer Science, MIT*, 2005.
5. J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MOBICOM*, 2005.
6. T. ElBatt, S. Krishnamurthy, D. Connors, and S. Dao. Power management for throughput enhancement in wireless ad-hoc networks. In *IEEE ICC*, 2000.
7. J. Huang, R. Berry, and M. Honig. Distributed interference compensation for wireless networks. In *IEEE JSAC, Vol. 24, No.5*, May 2006.
8. S. Kandukuri and S. Boyd. Optimal power control in interference-limited fading wireless channels with outage-probability specifications. In *IEEE Transactions of Wireless Communications, Vol. 1, issue 1, ISSN: 1536-1276, pp. 46-55*, Jan. 2005.
9. MadWifi-Driver. <http://madwifi.org>.
10. A. Muqattash and M. Krunz. Powmac: A single-channel power-control protocol for throughput enhancement in wireless ad hoc networks. In *IEEE JSAC, Vol. 23, No. 5, pp. 1067-1084*, May 2005.
11. T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan. Achieving mac layer fairness in wireless packet networks. In *MOBICOM*, 2000.
12. NLANR/iperf-version2.0.2. <http://dast.nlanr.net/projects/iperf/>.

13. J. Padhye, S. Agarwal, V. N. Padmanabhan, L. Qiu, A., and B. Zill. Estimation of link interference in static multi-hop wireless networks. In *IMC*, 2005.
14. Soekris/net4826. <http://www.soekris.com/net4826.htm>.
15. D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of the effects of transmission power control and blacklisting in wireless sensor networks. In *IEEE SECON*, 2004.
16. R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang. Distributed topology control for power efficient operation in multihop wireless ad hoc networks. In *INFOCOM*, 2001.
17. X. Yang and N. Vaidya. Priority scheduling in wireless ad hoc networks. In *MOBIHOC*, 2002.
18. H. Zhang, A. Arora, and P. Sinha. Learn on the fly: Beacon-free link estimation and routing in sensor network backbones. In *IEEE INFOCOM*, 2006.