

Self-management in chaotic wireless deployments

Aditya Akella · Glenn Judd · Srinivasan Seshan ·
Peter Steenkiste

Published online: 23 October 2006
© Springer Science + Business Media, LLC 2006

Abstract Over the past few years, wireless networking technologies have made vast forays into our daily lives. Today, one can find 802.11 hardware and other personal wireless technology employed at homes, shopping malls, coffee shops and airports. Present-day wireless network deployments bear two important properties: they are *unplanned*, with most access points (APs) deployed by users in a spontaneous manner, resulting in highly variable AP densities; and they are *unmanaged*, since manually configuring and managing a wireless network is very complicated. We refer to such wireless deployments as being *chaotic*.

In this paper, we present a study of the impact of interference in chaotic 802.11 deployments on end-client performance. First, using large-scale measurement data from several cities, we show that it is not uncommon to have tens of APs deployed in close proximity of each other. Moreover, most APs are not configured to minimize interference with their neighbors. We then perform trace-driven simulations to show that the performance of end-clients could

suffer significantly in chaotic deployments. We argue that end-client experience could be significantly improved by making chaotic wireless networks *self-managing*. We design and evaluate automated power control and rate adaptation algorithms to minimize interference among neighboring APs, while ensuring robust end-client performance.

Keywords Dense deployment · 802.11 access points · Interference · Power control · Channel assignment

1 Introduction

Wireless networking technology is ideal for many environments, including homes, airports, and shopping malls because it is inexpensive, easy to install (no wires), and supports mobile users. As such, we have seen a sharp increase in the use of wireless over the past few years. However, using wireless technology effectively is surprisingly difficult. First, wireless links are susceptible to degradation (e.g., attenuation and fading) and interference, both of which can result in poor and unpredictable performance. Second, since wireless deployments must share the relatively scarce spectrum that is available for public use, they often interfere with each other. These factors become especially challenging in deployments where wireless devices such as access points (APs) are placed in very close proximity.

In the past, most dense deployments of wireless networks were in campus-like environments, where experts could carefully manage interference by planning cell layout, sometimes using special tools [10]. However, the rapid deployment of cheap 802.11 hardware and other personal wireless technology (2.4 GHz cordless phones, bluetooth devices, etc.) is quickly changing the wireless landscape. Market estimates indicate that approximately 4.5 million WiFi APs were sold during the 3rd quarter of 2004 alone [13] and that the sales of

This work was supported by the Army Research Office under grant number DAAD19-02-1-0389, and by the NSF under grant numbers ANI-0092678, CCR-0205266, and CNS-0434824, as well as by IBM and Intel.

A. Akella (✉) · G. Judd · S. Seshan · P. Steenkiste
Computer Science Department, Carnegie Mellon University,
5000 Forbes Avenue, Pittsburgh, PA 15213, USA
e-mail: aditya@cs.cmu.edu

G. Judd
e-mail: glennj@cs.cmu.edu

S. Seshan
e-mail: srini+@cs.cmu.edu

P. Steenkiste
e-mail: prs@cs.cmu.edu

WiFi equipment will triple by 2009 [7]. The resulting dense deployment of wireless networking equipment in areas such as neighborhoods, shopping malls, and apartment buildings differs from past dense campus-like deployments in two important ways:

- **Unplanned.** While campus deployments are carefully planned to optimize coverage and minimize cell overlap, many recent deployments result from individuals or independent organizations setting up one or a small number of APs. Such spontaneous deployment results in variable densities of wireless nodes and APs and, in some cases, the densities can become very high (e.g. urban environments, apartments). Moreover, 802.11 nodes have to share the spectrum with other networking technologies (e.g., Bluetooth) and devices (e.g., cordless phones).
- **Unmanaged.** Configuring and managing wireless networks is difficult for most people. Management issues include choosing relatively simple parameters such as SSID and channel, and more complex questions such as number and placement of APs, and power control. Other aspects of management include troubleshooting, adapting to changes in the environment and traffic load, and making the wireless network secure.

We use the term *chaotic deployments or chaotic networks* to refer to a collection of wireless networks with the above properties. Such deployments provide many unique opportunities. For example, they may enable new techniques to determine location [15] or can provide near ubiquitous wireless connectivity. However, they also create numerous challenges. As wireless networks become more common and more densely packed, more of these chaotic deployments will suffer from serious contention, poor performance, and security problems. This will hinder the deployment and use of these infrastructures, negating many of the benefits offered by wireless networks.

The main goal of this paper is to show that interference in chaotic 802.11 deployments can significantly affect end-user performance. To this end, we first use large-scale measurements of 802.11 APs deployed in several US cities, to quantify current density of deployment, as well as configuration characteristics, of 802.11 hardware. Our analysis of the data shows that regions with tens of APs deployed in close proximity of each other already exist in most major cities. Also, most 802.11 users employ default, factory-set configurations for key parameters such as the transmission channel. Interestingly, we find that relatively new wireless technology (e.g., 802.11g) gets deployed very quickly.

We then simulate the measured deployment and configuration patterns to study the impact that unplanned AP deployments have on end-user performance. While it is true that the impact on end-user performance depends on the workloads imposed by users on their network, we do find that even when

the APs in an unplanned deployment are carefully configured to use the optimal static channel assignment, users may experience significant performance degradation, e.g. by as much of a factor of 3 in throughput. This effect is especially pronounced when AP density (and associated client density) is high and the traffic load is heavy.

To improve end-user performance in chaotic deployments, we explore the use of algorithms that automatically manage the transmission power and transmissions rates of APs and clients. In combination with careful channel assignment, our power control algorithms attempt to minimize the interference between neighboring APs by reducing transmission power on individual APs when possible. The strawman power control algorithm we develop, called Power-controlled Estimated Rate Fallback (PERF), reduces transmission power as long as the link between an AP and client can maintain the maximum possible speed (11 Mbps for 802.11b). Experiments with an implementation of PERF show that it can significantly improve the performance observed by clients of APs that are close to each other. For example, we show that a highly utilized AP-client pair near another such pair can see its throughput increase from 0.15 Mbps to 3.5 Mbps. In general, we use the term *self management* to refer to unilateral automatic configuration of key access point properties, such as transmission power and channel. Incorporating mechanisms for self-management into future wireless devices could go a long way toward improving end-user performance in chaotic networks.

The rest of the paper is structured as follows. We present related work in Section 2. In Section 3 we characterize the density and usage of 802.11 hardware across various US cities. Section 4 presents a simulation study of the effect of dense unmanaged 802.11 deployments on end-user performance. We present an analysis of power control in two-dimensional grid-like deployment in Section 5. In Section 6, we outline the challenges involved in making chaotic deployments self-managing. We describe our implementation of rate adaptation and power management techniques in Section 7. Section 8 presents an experimental evaluation of these techniques. We discuss other possible power control algorithms in Section 9 and conclude the paper in Section 10.

2 Related work

We first discuss current efforts to map 802.11 deployments. Then, we present an overview of commercial services and products for managing networks in general, and wireless networks in particular. Finally, we contrast our proposal for wireless self management (i.e., transmission power control and multi-rate adaptation) with related past approaches.

Several Internet Web sites provide street-level maps of WiFi hot-spots in various cities. Popular examples include

Wifi Maps [33], Wi-Fi-Zones.com [32] and JIWire.com [16]. Several vendors also market products targeted at locating wireless networks while on the go (see for example, Intego WiFi Locator [14]). Among research studies, the Intel Place Lab project [4, 15] maintains a database of up to 30,000 802.11b APs from several US cities. In this paper, we use hot-spot data from Wifi Maps.com, as well as the Intel Place Lab database of APs, to infer deployment and usage characteristics of 802.11 hardware. To the best of our knowledge, ours is the first research study to quantify these characteristics. We describe our data sets in greater detail in Section 3.

The general problem of automatically managing and configuring devices has been well-studied in the wired networking domain. While many solutions exist [30, 35] and have been widely deployed [9], a number of interesting research problems in simplifying network management still remain (e.g., [6, 26]). Our work in this paper compliments these results by extending them to the wireless domain.

In the wireless domain, several commercial vendors market automated network management software for APs. Examples include Propagate Networks' Autocell [22], Strix Systems' Access/One Network [29] and Alcatel OmniAccess' Air-View Software [3]. At a high-level, these products aim to detect interference and adapt to it by altering the transmit power levels on the access points. Some of them (e.g., Access/One) have additional support for managing load and coverage (or "coverage hole management") across multiple APs deployed throughout an enterprise network. However, most of these products are tailor-made for specific hardware (for example, AirView comes embedded in all Alcatel OmniAccess hardware) and little is known about the (proprietary) designs of these products. Also, these products are targeted primarily at large deployments with several tens of clients accessing and sharing a wireless network.

Also, several rate adaptation mechanisms that leverage the multiple rates supported by 802.11 have recently been proposed. For example, Sadeghi et al. [27] study new multi-rate adaptation algorithms to improve throughput performance in ad hoc networks. While our rate control algorithms, in contrast, are designed specifically to work well in conjunction with power control, it is possible to extend past algorithms such as [27] to support power control.

Similarly, traffic scheduling algorithms have been proposed to optimize battery power in sensor networks, as well

as 802.11 networks (see [19, 23]). In contrast, our focus in this paper is not on saving energy, *per se*. Instead we develop power control algorithms that enable efficient use of the wireless spectrum in dense wireless networks.

In general, ad hoc networks have recently received a great deal of attention and the issues of power and rate control have been also studied in the context of ad hoc routing protocols, e.g. [8, 11, 18, 28]. There are, however, significant differences between ad hoc networks and chaotic networks. First, ad hoc networks are multi-hop while our focus is on AP-based infrastructure networks. Moreover, nodes in ad hoc networks are often power limited and mobile. In contrast, the nodes in chaotic networks will typically have limited mobility and sufficient power. Finally, most ad hoc networks consist of nodes that are willing to cooperate. In contrast, chaotic networks involve nodes from many organizations, which are competing for bandwidth and spectrum. As we will see in Section 6, this has a significant impact on the design of power and rate control algorithms.

3 Characterizing current 802.11 deployments

To better understand the problems created by chaotic deployments, we collect and analyze data about 802.11 AP deployment in a set of metropolitan areas. We present preliminary observations of the density of APs in these metropolitan areas, and typical usage characteristics, such as the channels used for transmission and common vendor types.

3.1 Measurement data sets

We use four separate measurement data sets to quantify the deployment density and usage of APs in various U.S. cities. The characteristics of the data sets are outlined in Table 1. A brief description of the data sets follows:

1. **Place Lab:** This data set contains a list of 802.11b APs located in various US cities, along with their GPS coordinates. The data was collected as part of Intel's Place Lab project [4, 15] in June 2004. The Place Lab software allows commodity hardware clients like notebooks, PDAs and cell phones to locate themselves by listening for radio

Table 1 Characteristics of the data sets

Data set	Collected on	No. of APs	Stats collected per AP
Place Lab	Jun 2004	28475	MAC, ESSID, GPS coordinates
Wifi Maps	Aug 2004	302934	MAC, ESSID, Channel
Pittsburgh Wardrive A	Jul 2004	667	MAC, ESSID, Channel supported rates, GPS coordinates
Pittsburgh Wardrive B	Nov 2005	4645	MAC, ESSID, Channel supported rates, GPS coordinates, encryption

Table 2 Statistics for APs measured in 6 US cities (Place Lab data set)

City	Number of APs	Max AP degree (i.e., # neighbors)	Max. connected component size	No. of connected components
Chicago	2370	20	54	369
Washington D.C.	2177	39	226	162
Boston	2551	42	168	320
Portland	8683	85	1405	971
San Diego	7934	54	93	1345
San Francisco	3037	76	409	186

beacons such as 802.11 APs, GSM cell phone towers, and fixed Bluetooth devices.

- WifiMaps:** Wifi Maps.com [33] provides an online GIS visualization tool, to map wardriving results uploaded by independent users onto street-level data from the US Census. We obtained access to the complete database of wardriving data maintained at this website as of August 2004. For each AP, the database provides the AP's geographic coordinates, zip code, its wireless network ID (ESSID), channel(s) employed and the MAC address.
- Pittsburgh wardrive A:** This data set was collected on July 29, 2004, as part of a small-scale wardriving effort which covered portions of a few densely populated residential areas of Pittsburgh. For each unique AP measured, we again collected the GPS coordinates, the ESSID, the MAC address and the channel employed.
- Pittsburgh wardrive B:** This data was collected on November 22, 2005. This wardrive was more extensive and thorough than the Pittsburgh wardrive A dataset, and as a result, contains far more access points. In addition to the information collected in the A dataset, we also collected information regarding the use of encryption, the exact data rates offered by APs, and the physical areas covered by APs.

3.2 Measurement observations

In this section, we analyze our data sets to identify real-world deployment properties that are relevant to the efficient functioning of wireless networks. The reader should note that data analyzed here provides a gross underestimate of any real-world efficiency problem. First, none of above data sets are complete—they may fail to identify many APs that are present and they certainly do not identify non-802.11 devices that share the same spectrum. Second, the density of wireless devices is increasing at a rapid rate, so contention in chaotic deployments will certainly increase dramatically as well. Because of these properties, we believe these data sets will lead us to underestimate deployment density. However, these data sets are not biased in any specific way and we expect our other results (e.g. channel usage, AP vendor and 802.11 g deployment) to be accurate.

3.2.1 802.11 Deployment density

First, we use the location information in the Place Lab data set to identify how many APs are within interference range of each other. For this analysis, we conservatively set the interference range to 50 m, which is considered typical of indoor deployments. We assume two nodes to be “neighbors” if they are within each other's interference range. We then use this neighborhood relationship to construct “interference graphs” in various cities.

The results for the analysis of the interference graphs in six US cities are shown in Table 2. On average we note 2400 APs in each city from the Place Lab dataset. The third column of Table 2 identifies the maximum degree of any AP in the six cities (where the degree of an AP is the number of other APs in interfering range). In San Francisco, for example, a particular wireless AP suffers interference from about 80 other APs deployed in close proximity.

In Fig. 1, we plot a distribution of the degrees of APs measured in the Place Lab data set. In most cities, we find several hundreds of APs with a degree of at least 3. In Portland, for example, we found that more than half of the 8683 nodes measured had 3 or more neighbors. Since only three of the 802.11b channels are non-overlapping (channel 1, 6 and 11), these nodes will interfere with at least one other node in their vicinity.

The fourth column in Table 2 shows the size of the maximum connected component in the interference graph of a


**Fig. 1** Distribution of AP degrees (Place Lab data set)

Table 3 Statistics for Pittsburgh APs (Pittsburgh wardrive B data set)

City	Number of APs	Max AP degree (i.e., # neighbors)	Max. connected component size	No. of connected components
Pittsburgh	4645	48	853	8

city. The final column shows the number of connected components in the interference graph. From these statistics, we find several large groups of APs deployed in close proximity. Together, these statistics show that dense deployments of 802.11 hardware have already begun to appear in urban settings. As mentioned earlier, we expect the density to continue to increase rapidly. Our analysis of the interference graphs in Pittsburgh based on the wardrive B dataset is shown in Table 3 and confirms the above observations regarding density.

3.2.2 802.11 Usage: Channels

Table 4 presents the distribution of channels used by APs in the Pittsburgh wardrive A and B data sets. This provides an indication of whether users of APs manage their networks at all. Notice that many APs transmit on channel 6, the default on many APs, and only a third use the remaining two non-overlapping channels in 802.11b (i.e., channels 1 and 11). While this does not identify particular conflicts, this distribution suggests that many of the APs that overlap in coverage are probably not configured to minimize interference.

3.2.3 802.11 Variants

The Pittsburgh wardrive A data set contains information about rates supported for about 71% of the measured APs, or 472 out of the 667. We use this information to classify these APs as 802.11b or 802.11g. We find that 20% of the classified APs, or 93, are 802.11g. Given the relatively recent standardization of 802.11g (June 2003), these measurements

Table 4 Channels employed by APs in the Pittsburgh A and B data sets

Channel	%-age of APs (wardrive A)	%-age of APs (wardrive B)
1	15.55	13.14
2	0.86	1.12
3	2.37	1.45
4	0.86	0.99
5	0.65	0.69
6	50.97	52.48
7	1.73	1.15
8	0.43	1.32
9	1.30	2.03
10	4.32	3.18
11	20.95	22.45

suggest that new wireless technology gets deployed relatively quickly. Table 6 repeats this analysis for the more recent wardrive B data. We find that 802.11 g continues to be adopted at a rapid rate, and is approximately 43%. 802.11 b variants make up just over half of the market while legacy 802.11 systems are nearly non-existent.

3.2.4 Supported rates

Even two networks that conform to the same standard can operate in a very different manner. One example of this is seen by examining the transmission rates supported by deployed access points. Figure 5 shows the supported rates advertised by the access points observed in the Pittsburgh B wardrive. The 802.11b and b+ networks support 6 different sets of rates while the g networks support 15 different sets of rates. While many of the rate sets are uncommon, it is clear that one cannot assume that all rates of a given standard are available.

3.2.5 Vendors and AP management support

To determine popular AP brands, we look up the MAC addresses available in the Wifi Maps data set against the IEEE Company_id assignments [12] to classify each AP according to the vendor. For the APs that could be classified in this manner (2% of the APs in the Wifi Maps data set did not have a matching vendor name), the distribution of the vendors is shown in Table 7. Notice that Cisco products (Linksys and Aironet) make up nearly half of the market. This observation suggests that if future products from this vendor incorporated built-in mechanisms for self-management of wireless networks this could significantly limit the impact of interference in chaotic deployments.

To understand if specific models incorporate software for configuration and management of wireless networks, we survey the popular APs marketed by the top three vendors in Table 7. All products (irrespective of the vendors) come with software to allow users to configure basic parameters for their wireless networks, such as ESSID, channel and security settings. Most “low-end” APs (e.g., those targeted for deployment by individual home users) do not include any software for *automatic* configuration and management of the wireless network. Some of the products targeted at enterprise and campus-style deployments, such as Cisco Aironet 350 series, allow more sophisticated, centralized management of parameters such as transmit power levels, selecting

non-overlapping channels, etc. across several deployed APs. Since these products are targeted at campuses, they are too expensive for use in smaller-scale deployments such as apartment-buildings.

3.2.6 Security settings

Table 8 displays security settings gleaned from the wardrive B data for two parameters: SSID visibility and encryption. Encrypted networks encrypt the contents of data packets in order to hide them from eavesdroppers. As an additional measure of security, some networks hide their SSID in order to discourage unauthorized network access.

The access points in the Pittsburgh B data set are nearly evenly split between open networks and encrypted networks. Roughly 12% of networks hide their SSID to discourage unauthorized access.

3.2.7 Coverage area

The amount of interference in chaotic networks is largely determined by access point density and the coverage area of access points. To gain insight into access point coverage in chaotic networks, we computed approximate lower bounds on the areas covered by each access point in the Pittsburgh wardrive B dataset (our calculation neglects potential “holes” in coverage). We took the observed measurement locations of each access point, computed a convex hull between these points (this gives us a rough estimate of the coverage area to the extent we can measure it), and then computed the area of each convex hull. Note that the area estimated for each AP in this manner is, in fact, a lower bound on the actual area covered by each access point. Figure 2 plots the estimated coverage areas for all access points measured in Pittsburgh. Access points with estimated coverage areas less than 400 square meters are omitted.




Fig. 2 AP coverage

Clearly there is a large variation in access point coverage with the largest area in excess of two square kilometers. This variation comes from three factors: measurement error, transmission power, and physical variation in the RF propagation environment. Of particular interest are the few access points that have extremely large coverage areas. In this dataset, the area covered by an access point had little to do with the network type (i.e., b, g or b+) despite claims frequently used for marketing purposes. In fact, the single legacy 802.11 system we observed in this data set has the fifth largest estimated coverage area. We found that, as expected, the physical environment is a much more dominant factor in determining the coverage area than the type of network used.

Figure 3 shows the area covered by the Pittsburgh B wardrive. Each pin on the map depicts a measurement location. The oval shows the area where the access points with the largest estimated coverage areas were observed. These were all located near the Monongahela river in Pittsburgh where there is ample open space for free-space radio propagation.

While such large coverage areas are convenient for lowering the amount of equipment required to cover large amounts of territory, they present challenges in chaotic networks, as they greatly increase the number of potential interferers. In other words, chaotic networks may actually benefit from walls, trees, and other obstacles to RF propagation, as these obstacles limit the scope of interference. Large open areas require mechanisms such as transmit power control in order to mitigate the effects of interference.

3.2.8 An aside: anomalous operation

We have observed two cases of anomalous network operation in the Pittsburgh wardrive B data set. The first is a minor issue of 5 access points (all Linksys) with rate sets containing duplicate rate advertisements as seen in line 5 of Table 5.

The second more serious issue is that several Netgear access points appear to have been given the same MAC address (the MAC addresses have the form 00:90:4c:7e:00:xx). We observe three distinct MAC addresses that have this problem. In each case several access points share the MAC in question. That the access points are distinct is revealed by the fact that we observe multiple SSIDs, distinct channels, and an unreasonably large coverage area for the given MACs. These sets have at least 4, 6, and 7 access points that share the MACs; the true number is difficult to determine due to the fact that the default SSID could be shared by multiple distinct MACs. We found an additional MAC in the Wifi Maps database that is shared by access points in several distinct locations across the United States.

Table 5 Supported rate sets in the Pittsburgh B data set

Network Type	Rate Set	Number	Percentage
b	[11.0]	8	0.17
b	[5.5 11.0]	33	0.71
b	[2.0 5.5 11.0]	2	0.04
b	[1.0 2.0 5.5 11.0]	1977	42.70
b	[1.0 5.5 11.0 11.0] (sic)	5	0.11
b+	[1.0 2.0 5.5 11.0 22.0]	578	12.48
g	[5.5 11.0 54.0]	1	0.02
g	[11.0 36.0 48.0 54.0]	56	1.21
g	[1.0 2.0 5.5 6.0 9.0 11.0]	1	0.02
g	[11.0 24.0 36.0 48.0 54.0]	1	0.02
g	[5.5 11.0 24.0 36.0 48.0 54.0]	1	0.02
g	[1.0 2.0 5.5 11.0 6.0 9.0 12.0 18.0]	71	1.53
g	[1.0 2.0 5.5 6.0 9.0 11.0 12.0 18.0]	280	6.05
g	[11.0 12.0 18.0 24.0 36.0 48.0 54.0]	1	0.02
g	[1.0 2.0 5.5 11.0 6.0 12.0 24.0 36.0]	318	6.87
g	[1.0 2.0 5.5 11.0 6.0 12.0 24.0 54.0]	2	0.04
g	[1.0 2.0 5.5 11.0 18.0 24.0 36.0 54.0]	1272	27.47
g	[5.5 6.0 9.0 11.0 12.0 18.0 24.0 36.0]	1	0.02
g	[6.0 9.0 11.0 12.0 18.0 24.0 36.0 48.0]	1	0.02
g	[1.0 2.0 5.5 11.0 6.0 9.0 12.0 18.0 24.0 36.0 48.0 54.0]	4	0.09
g	[1.0 2.0 5.5 6.0 9.0 11.0 12.0 18.0 24.0 36.0 48.0 54.0]	16	0.35
legacy 802.11	[1.0 2.0]	1	0.02

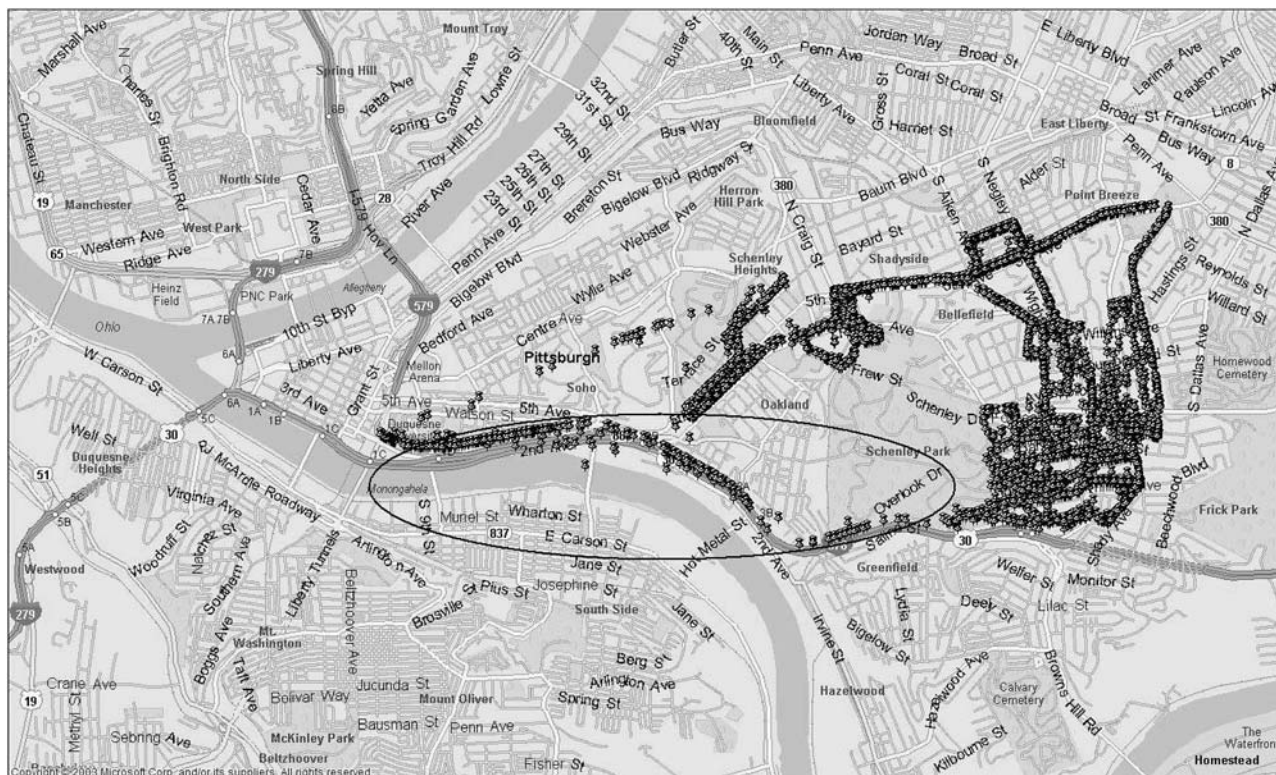


Fig. 3 Pittsburgh B wardrive Map

Table 6 Networks types observed in the Pittsburgh wardrive B data set

Network type	Number	Percentage
b	2025	43.73
b+	578	12.48
g	2026	43.75
legacy 802.11	1	0.02

Table 7 Popular AP vendors (Wifi Maps data set)

Vendor	Percentage of APs
Linksys (Cisco)	33.5
Aironet (Cisco)	12.2
Agere systems	9.6
D-Link	4.9
Apple computer	4.6
Netgear	4.4
ANI communications	4.3
Delta networks	3.0
Lucent	2.5
Acer	2.3
Others	16.7
Unclassified	2

Table 8 AP security settings

	Broadcast SSID	Hide SSID	Total
Encrypted	39.87	9.65	49.52
Unencrypted	48.09	2.26	50.35
Unknown	0.11	0.02	0.13
Total	88.07	11.93	100

4 Impact on end-user performance

In order to quantify the impact of the deployment and usage characteristics of 802.11b APs on the Internet performance observed by end-users, we conducted trace-driven simulations using the publicly available GloMoSim simulator [34]. We simulated two real deployment topologies shown in Figs. 4(a) and (b). These topologies, which we refer to as \mathcal{R} and \mathcal{S} , were collected during the Pittsburgh wardrive. They

Table 9 Table (a) shows minimum required SNR for Prism 2.5. Table (b) shows maximum 802.11 b throughput

Rate (Mbps)	Minimum SNR (dB)	Rate (Mbps)	Throughput (Mbps)
1	3	1	0.85
2	4	2	1.7
5.5	8	5.5	3.5
11	12	11	4.9
(a)		(b)	

contain 20 and 29 APs respectively. We use the following settings and assumptions in our simulations: (1) Each node in the map corresponds to an AP; (2) Each AP has D clients (e.g., laptops) associated with it. We vary D between 1 and 31 (3) Clients are located less than 1m away from their respective APs and do not move; (4) Unless otherwise specified, we assume that all APs transmit on channel 6; (5) All APs employ a fixed transmit power level of 15 dBm, unless otherwise specified (This is the default setting in most commercial APs); (6) All APs transmit at a single rate, 2 Mbps (there is no multi rate support in GloMoSim). At these settings, the transmission and interference ranges are 31 m and 65 m, respectively; (7) RTS/CTS is turned off. This is the default setting in most commercial APs; And, (8) We use a modified two-ray path loss model for large-scale path loss, and a Ricean fading model with a K-factor of 0 for small scale fading [25].

Intuition suggests that the impact of interference in chaotic wireless deployments depends, to a large extent, on the workloads imposed by users. If most APs are involved in just occasional transmission of data to their users, then it is very likely that users will experience no degradation in performance due to interference from nearby APs. A key goal of our simulations, then, is to systematically quantify the precise impact of user workloads on eventual user performance. To achieve this, we simulate two types of user workloads over the above simulation set-up. These workloads differ mainly in their relative proportions of HTTP (representing Web-browsing activity) and FTP (representing large file downloads) traffic.

In the first set of workloads, called `http`, we assume that the clients are running HTTP sessions across their APs. The HTTP file size distribution is based on a well-known model for HTTP traffic [20]. On a client, each HTTP transfer is separated from the previous one by a think time drawn from a Poisson distribution with a mean of s seconds. We vary s between the values of 5 s and 20 s (We also simulated HTTP workloads with 10 s, 30 s and 60 s sleep times. The results are qualitatively similar and are omitted for brevity). The average load offered by the HTTP client is 83.3 Kbps for a 5 s sleep time, and 24.5 Kbps for a 20 s sleep time. There is no other interfering traffic in the `http` workload.

The second set of workloads, called `comb-ftpi`, is similar to the `http` workload with the exception of i clients in the entire set-up running long-lived FTP flows for the duration of the simulation. We vary i between 1 and 3. The average load offered by the FTP clients in our simulation is 0.89 Mbps. We simulate either set of workloads for 300 s.

4.1 Interference at low client-densities and traffic volumes

First, we conduct simulations with the `http` and `comb-ftp1` workloads, and low client densities ($D = 1$). All the




Fig. 4 Two simulation topologies derived from the Pittsburgh wardrive data set; the x and y axes are in meters

results we present in this section are for simulations over the \mathcal{R} topology (Fig. 4(a)), unless otherwise specified. In general, we note that the experimental results for the \mathcal{S} topology are identical.

The results for the \mathcal{R} topology are shown in Fig. 5. The performance measurements are the average of 5 different simulation runs; the variance between runs is not shown since it was low. The x-axis in these pictures is the “stretch” parameter which allows us to tune the density of APs per square meter in a given simulation. A simulation with a stretch of l indicates that the distance between a pair of APs in the simulation topology is l times the actual distance in the original topology. The distance between an AP and its clients does not change. The higher the value of stretch, the lower the likelihood of interference between nodes in the simulation topology. For the \mathcal{R} topology, we note that at $stretch \approx 20$, the nodes are completely out of each others’ interference range. Also, in our simulations, beyond $stretch = 10$, we see little impact of interference between nodes on user performance. In either figure, the y-axis shows the average normalized performance of HTTP (Fig. (a)) or FTP flows (Fig. (b)) in our

simulations. Normalized HTTP (FTP) performance is simply the ratio of the average throughput of an HTTP (FTP) flow to the throughput achieved by an FTP bulk transfer when operating in isolation, i.e., 0.89 Mbps. This can be viewed as the amount of work a user completed during a fixed time interval, relative to the maximum achievable work.

Notice that, for workloads with an “aggressive” HTTP component (i.e., think time of 5 s), the performance of the HTTP flows improves until $stretch = 10$; beyond this point performance stays relatively flat. For less aggressive HTTP workloads (i.e., think interval of 20 s), the impact on the performance of the HTTP flows is less severe. The performance of the FTP flow in $comb-ftp_1$ workload is shown in Fig. 5(b). When the HTTP component of this workload is aggressive ($s = 5$ s), the performance of the lone FTP flow suffers by about 17%. With a not-so-aggressive HTTP component, as expected, the impact on the FTP flow is minimal.

So far, we studied the impact of interference under relatively “light-weight” user traffic at each access points. In the next two sections, we vary two important factors determining the client load—the density of clients per AP and

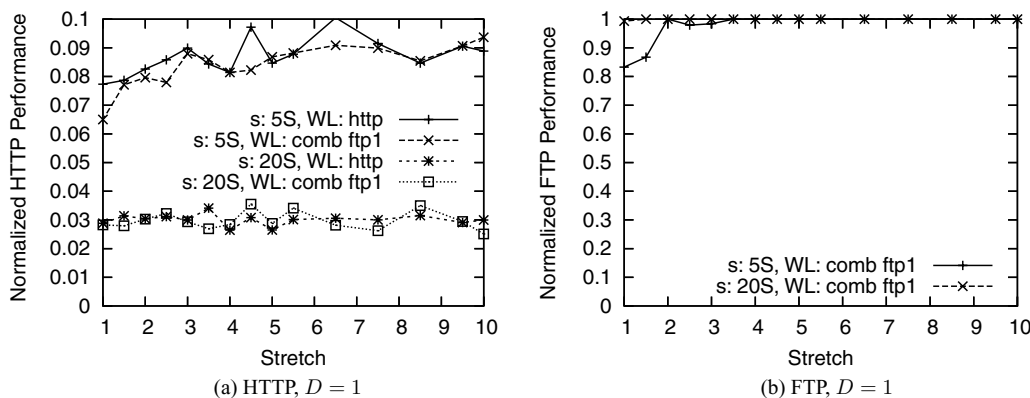


Fig. 5 Average performance of HTTP and FTP flows at low client densities ($D = 1$) and low levels of competing FTP traffic ($http$ and $comb-ftp_1$ workloads) for topology \mathcal{R}




Fig. 6 Average performance of HTTP and FTP flows at greater client densities ($D = 3$) for topology \mathcal{R}

the traffic volume of the clients—to create more aggressive interference settings.

4.2 Impact of client densities and traffic load

Impact of client density. Figures 6(a) and (b) show the average performance of the individual HTTP and FTP sessions, respectively, in the *comb-ftp₁* and *http* workloads, for a high number of clients associated per AP ($D = 3$). The performance in both cases suffers significantly under high client densities: From Fig. 6(a), HTTP performance is lowered by about 65% (compare *stretch* = 1 with *stretch* = 10) due to interference between aggressive HTTP flows ($s = 5$ s). The same is true for the performance of the FTP flow in Fig. 6(b). For a less aggressive HTTP component ($s = 20$ s) the performance of the HTTP flows is 20% inferior, while the FTP flow suffers by about 36%. For the \mathcal{S} topology, the performance impact on HTTP and FTP is 57% and 60% respectively, with an aggressive HTTP component in interfering traffic (results omitted for brevity).

Impact of traffic volume. Figures 7(a) and (b) show the average performance of the HTTP and FTP flows, respectively, in simulations with a few more competing FTP flows—

i.e., the *comb-ftp₂* and *comb-ftp₃* workloads—for $D = 3$. The performance impact on HTTP and FTP flows is slightly more pronounced, even for the cases where the HTTP component of these workloads is not very aggressive (see the curves corresponding to $s = 20$ s in Fig. 7(b)). Results for the \mathcal{S} topology are identical and omitted for brevity.

Using realistic channel assignments. We also performed simulations on the 20-node topology, where the APs were statically assigned channels based on the distribution in Table 4. However, we note similar levels of interference and impact on performance as observed above. This is because more than half the APs in this simulation were assigned channel 6, which was the most predominant channel employed by most APs according to our measurements.

4.3 Limiting the impact of interference

In this section, we explore the effect of two simple mechanisms on mitigating interference in chaotic networks: First, we study if an optimal static allocation of non-overlapping channels across APs could eliminate interference altogether. Second, we present a preliminary investigation of the effect

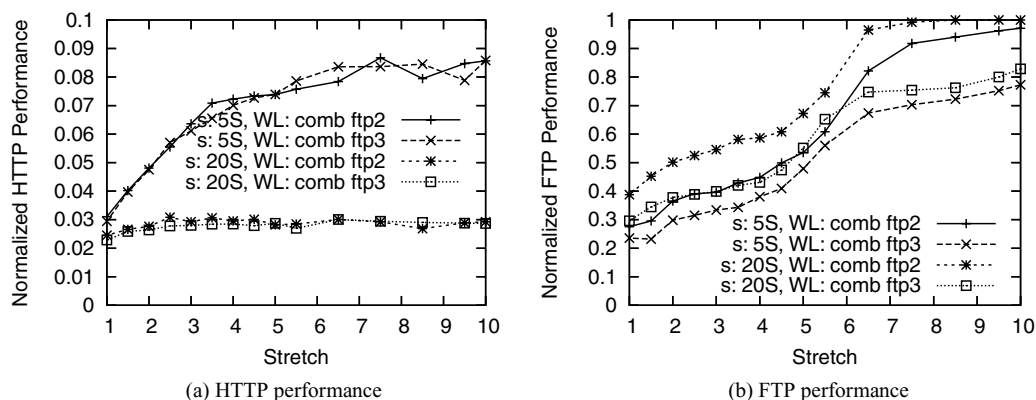


Fig. 7 Average performance of HTTP and FTP flows at higher competing FTP traffic levels (i.e., the *comb-ftp*{2, 3} workloads) and for $D = 3$ in topology \mathcal{R}




Fig. 8 Two sub-topologies obtained from static optimal channel allocation. Figure (a) shows an 8-node sub-topology of map \mathcal{R} in Fig. 4(a). All 11 nodes were assigned channel 1 by the optimal static allocation

algorithm. Similarly, Fig. (b) shows an 11-node sub-topology of map \mathcal{S} in Fig. 4(b)

of reducing the transmit power levels at APs on the interference experienced. We also investigate how transmit power control improves the total capacity of a chaotic, network, as well as the fairness in the allocation of the capacity among individual APs.

4.3.1 Effect of optimal static channel allocation

We performed simulations on the topologies of Fig. 4, where the APs are statically assigned one of the three non-overlapping channels (1, 6 and 11) such that no two neighboring APs share a channel, whenever possible. Figures 8(a) and (b) show the lay-out of APs on maps \mathcal{R} and \mathcal{S} that were all assigned channel 1 by this scheme. We only show results for the \mathcal{R}_{sub} sub-topology in Fig. 8(a). The results for \mathcal{S}_{sub} are identical and omitted for brevity.

The performance of HTTP and FTP flows in these simulations are shown in Figs. 9(a) and (b), respectively. The average performance of both HTTP and FTP flows improves significantly. Comparing with Figs. 7(a) and (b) respectively, we note that the performance curves “flatten out” earlier on

account on the sparse nature of the interference graph. Nevertheless, the impact of interference can still be seen: the average HTTP performance is about 25% inferior at $stretch = 1$ compared to the case when no nodes interferes with another ($stretch = 10$). FTP performance, similarly, is far from optimal. These observations suggest that while optimal static channel allocation reduces the impact of interference, it cannot eliminate it altogether.

4.3.2 Impact of transmit power control

We augment above simulations of optimal static channel allocation with more conservative (lower) power settings on the APs: we forced the APs in Figs. 8(a) and (b) to use a power level of 3 dBm. This yields a transmission range of 15 m, which is half the range from using the default power level of 15 dBm. Next, we show how this improves HTTP and FTP performance, as well as the total network capacity. We show results for the \mathcal{R}_{sub} topology next. The results for \mathcal{S}_{sub} are very similar.




Fig. 9 Performance of HTTP and FTP flows with optimal static assignment of APs to the three non-overlapping channels for topology \mathcal{R} . The transmit power level is set at 15 dBm, corresponding to a reception range of 31 m

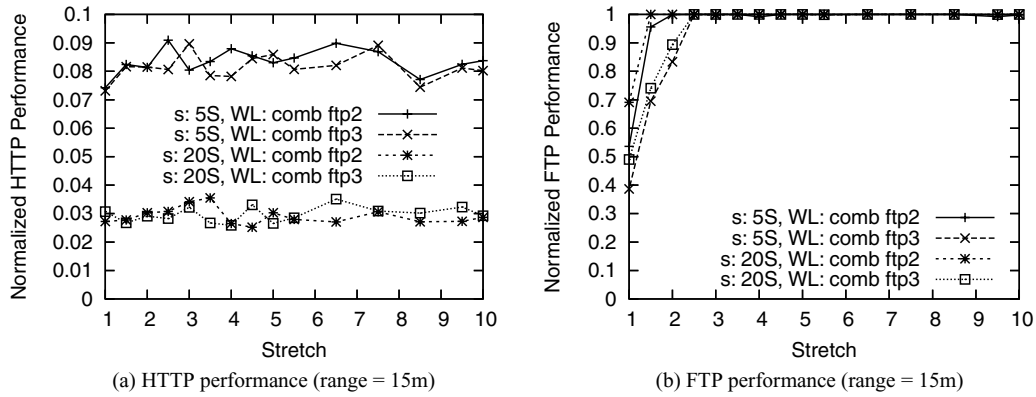


Fig. 10 Performance of HTTP and FTP flows with optimal static channel assignment of APs and the transmit power level set at 3 dBm. This corresponds to a reception range of 15 m

Improvement in application performance. The performance results for HTTP and FTP flows in these simulations are shown in Figs. 10(a) and (b) respectively. Compared with Figs. 9(a) and (b), the performance of individual flows improves significantly. The interference among nodes is lowered, as can be seen by both the performance curves flattening out at *stretch* = 2. These results show that transmit power control, in conjunction with a good channel allocation mechanism, could help reduce the impact of interference in chaotic networks substantially.

Improvement in network capacity. In Fig. 11 we show how transmit power control improves user performance for a workload composed fully of bulk FTP transfers (i.e., each AP has one user associated with it, and the AP runs an FTP bulk transfer to the user). This simulation sheds light on how careful management of APs can improve the total capacity of the network. When APs are completely unmanaged, the capacity of a densely packed network of APs is only 15% of the maximum capacity (see Fig. 11(a)). Static channel allocation of APs improves the capacity two-fold. Lowering

the transmit power on APs improves capacity by nearly an additional factor of 2.

Fairness. In Fig. 11(b), we show the fairness in the throughputs achieved by individual FTP flows to understand if the performance of certain APs in a chaotic deployment suffers significantly compared to others. Our fairness metric is derived from [5] and is defined as $\frac{(\sum x_i)^2}{n \sum x_i^2}$, where x_i 's are the throughputs of individual flows.

For the highest densities of access points, we see that poor management results in unfair allocation of capacity across access points. Channel allocation coupled with transmit power control immediately ensures a highly equitable allocation: except for the highest density, the fairness of allocation is above 0.9.

In this section, we used simulations on two distinct sets of deployment topologies to study the impact of transmit power reduction on network performance. The key observation from our simulations is that end-user performance can suffer significantly in chaotic deployments, especially when the interference is from aggressive sources (such as

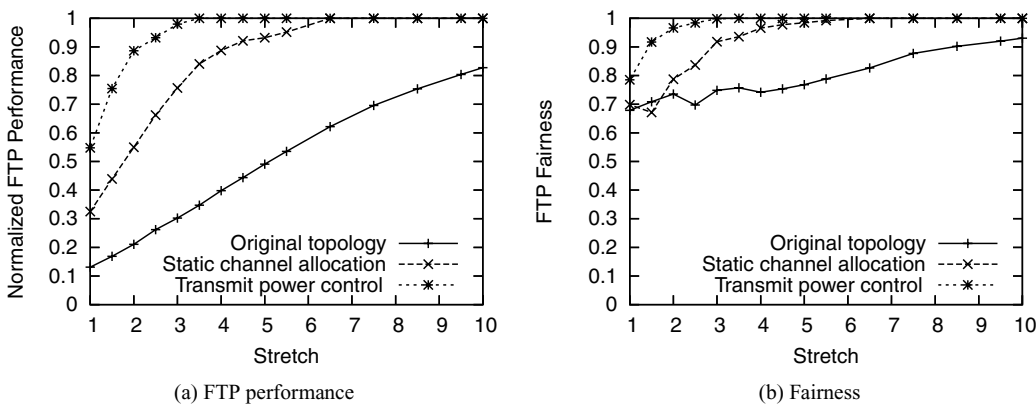


Fig. 11 Performance FTP flows with and without optimal channel assignment and AP transmit power control. The workload is composed of FTP flows between each client and its AP, with $D = 1$. Figure (a) shows

the performance of FTP flows in the simulations. Figure (b) shows the fairness index for the throughput achieved by the FTP flows

bulk FTP transfers). We showed that careful management of APs, via transmit power control (and static channel allocation), could mitigate the negative impact on performance. Moreover, transmit power control can also enable an equitable allocation of capacity among interfering APs. A key drawback of our simulations, however, is the lack of support for multi-rate adaptation. We further explore the benefits of transmit power control in conjunction with multi-rate adaptation in the next section.

5 Benefits of transmit power reduction

Next, we use a simple model of wireless communication applied to a 2-D grid topology as shown in Fig. 12 to quantify the advantages of transmit power control. We also model the impact of rate adaptation. For our analysis, we assume that each AP sends traffic to a single client at a fixed distance (d_{client}) from the AP. In practice, if this transfer used TCP, we would expect a small amount of traffic from the client to AP due to the TCP acknowledgements. As the amount of uplink traffic is small, we only consider the downlink traffic to simplify our analysis. We also ignore many real-world effects such as multipath fading and channel capture. We do not use these simplifying assumptions in subsequent sections. In particular, we stress that our algorithms are designed to operate in a symmetric fashion on both uplink and downlink traffic.

We examine a range of transmit power levels and traffic loads. For each transmit power level and traffic load pair, we determine the minimum physical spacing required between APs (d_{min}) to support the specified load. To compute this, we first calculate the medium utilization required

by each AP: $utilization_{AP} = load / throughput_{max}$. The maximum throughput is determined by first calculating path loss from the AP to the client (in dB) as: $pathloss = 40 + 3.5 * 10 * \log(d_{client})$; this is based on the pathloss model from [25] with constants that correspond to measurements collected in our local environment.

Received signal strength can then be computed as $RSS = txPower - pathloss$, and the signal-to-noise ratio is $SNR = RSS - noiseFloor$. We choose the noise floor to be -100 dBm, which is typical for our hardware. Using SNR and the data in Tables 9(a) and (b) (based on measurements presented in [1]) we then determine the maximum transmission rate that can be used for the AP-client link and the corresponding maximum throughput for that rate.

After we have computed the utilization for a single link, we determine the medium utilization at each AP by summing the utilization of all in-range APs. We determine that two APs are in range by computing the RSS between them using the same formula used above. We consider the candidate AP to be in range of the local AP if $RSS > interferenceThreshold$. We set $interferenceThreshold$ to -100 dBm for our calculations.

Figure 13 shows the results of our calculations for a client distance of 10 meters and loads ranging from 0.1 Mbps to 1.1 mbps. Other client distances and loss parameters would shift or scale the graph, but the trends would remain the same. The 2, 5.5, and 11 Mbps regions shown near the bottom of the chart specify the maximum transmission rate that can be used for each power level. Consider a specific point on this graph, such as the point at -15 dBm power on the x -axis; using the 1.1 Mbps load line, this translates into a 20 m AP distance on the y -axis. This implies that if you want to transmit at 1.1 Mbps from each AP at -15 dBm, the APs must be at least 20 m apart in order to avoid overloading the wireless link. In addition, the solid-black line parallel to the x -axis indicates that each AP uses a transmission rate of 5.5 Mbps to communicate with its client. For the simulations in the




Fig. 12 Computing minimum AP spacing for a grid topology




Fig. 13 Minimum AP distance vs. Tx power ($d_{client} = 10 m$)

previous section we typically fixed the transmit power and then increased the stretch; this corresponds to picking a point on the x -axis and moving up vertically.

We draw a few key conclusions from Fig. 13. Clearly the minimum distance between APs that can be supported decreases (i.e. maximum supported density increases) dramatically as the transmission power (in dBm) is decreased. We also find that high AP density and higher loads require transmit power levels below 0 dBm. This is the lowest transmit power available from commercial hardware that we are aware of. Adding support for lower transmit power levels to wireless hardware would be a simple way of improving the density of APs that can be supported. Secondly, the graph can also be used to determine the upper bound on the power level that should be employed (x -axis) in order to achieve a certain throughput, given a certain inter-AP distance (y -axis). Using a higher power level will typically not affect (i.e. not decrease or increase) the performance for that node, but will reduce performance for other, nearby nodes. This is the basis for one of the power control algorithms discussed in the next section. Finally we note that the highest densities require the use of very low transmission power, forcing nodes to use a transmission rate under 11 Mbps. This suggests that, when their traffic requirements are low, it may be advantageous if nodes voluntarily reduce not only their transmission power but also their transmission rate since it could increase the overall network capacity in very dense networks. We will revisit this issue in Section 9.

6 Deployment challenges

Power control offers a simple but powerful technique for reducing interference. The tradeoffs are obvious: reducing the power on a channel can improve performance for other channels by reducing interference, but it can reduce the throughput of the channel by forcing the transmitter to use a lower rate to deal with the reduced signal-to-noise ratio. As a result, we must carefully consider the incentives that users may have for using such techniques. In practice, the incentives for using power control are complex and we have to distinguish between the techniques that are applicable to campus deployments and chaotic wireless networks.

In campus environments, there are a number of APs under the control of a single organization. This organization is in a position to do power control in each cell in a way that optimizes some global network metric, e.g. total network throughput or fairness. An additional important consideration is that in campus networks a user can obtain service from any of the APs in transmission range. Therefore, any design may need to carefully consider issues such as load-balancing of users across APs along with power control.

In chaotic networks, the infrastructure is controlled by multiple organizations, and, unfortunately, their priorities often conflict. For example, for a home network consisting of a single AP, the best strategy is to always transmit at maximum power, and there is no incentive to reduce power and, thus, interference. The results in the previous section show that such a “Max Power” strategy, when employed by multiple APs, will result in suboptimal network performance. This implies that while a single node can improve its performance by increasing power, it can actually obtain better performance if it, and all of its neighbors, act socially and reduce their transmission power appropriately. This is analogous to the tradeoffs between selfish and social congestion control in the Internet [2]. While a node can improve performance by transmitting more quickly in the Internet, this can result in congestion collapse and degraded performance for all. We believe that similar factors that drove the wide deployment of congestion control algorithms will drive the deployment of power control algorithms. We should note that an added side incentive for the deployment of automatic power control is that it limits the propagation of an AP’s transmission which, in turn, limits the opportunity of malicious users eavesdropping on any transmission.

Our work focuses on socially responsible power control algorithms that would work well in chaotic environments. We call such power control algorithms “socially responsible” to differentiate them from approaches that require global coordination across multiple access points (e.g., for campus-wide wireless networks). Our algorithms are targeted at individual access points and clients, which behave in an altruistic manner, agnostic to the actions of other APs and clients. Our algorithms could also work in campus scenarios. However, we do not consider issues such as AP load-balancing which arise in such environments. Extending our design to campus deployments is left for future work.

Note that while our algorithms are targeted at nodes behaving in an altruistic manner, there are also practical considerations that make them more feasible than simply relying on the altruism of end users would. In particular, these algorithms are implemented not by end users, but by equipment vendors. From an equipment vendor’s point-of-view, reducing interference is beneficial. Moreover, regulatory mandates already limit transmit power, and could be extended to require dynamic adjustment of transmit power in order to increase spatial reuse and potentially allow for higher transmit power limits which would clearly benefit both end users and equipment vendors. Finally, as discussed in Section 3.2, we find that new technology is quickly adopted in chaotic networks, and that many users in chaotic networks do not change factory default settings. Hence, vendor implemented intelligent transmit power control could be deployed relatively quickly and would be widely adopted.

7 Transmission power and rate selection

In order to characterize how power adaptation affects both network-wide and individual user throughput, we ran experiments with several rate selection algorithms implemented on both APs and clients. In this section, we describe the fixed-power rate selection algorithms and adaptive-power algorithms that we evaluate. Before we introduce the algorithms, we briefly describe our implementation environment for the rate selection algorithms.

7.1 Rate selection implementation

Our experiments use a NIC based on the Prism chipset running the 2.5 version of the firmware. The driver is a modified version of the HostAP [21] Prism driver for Linux, which was extensively modified to give fine-grained control of rate selection and transmission power to the driver.

The driver achieves per-packet control over transmission rate by tagging each packet with the rate at which it should be sent. The Prism 2.5 firmware will then ignore its internal rate selection algorithm and send the packet at the specified rate; all firmware retransmissions will use this same rate. The Prism 2.5 firmware also allows us to take control over the retransmission of packets (in the driver) by tagging each packet with a specified number of retries. For instance, setting the packet retry count to zero tells the firmware to attempt no retries in firmware, and to merely inform the driver if a packet is not acknowledged. This allows us to completely replace the firmware rate selection and retransmission algorithms. While this gives us the ability to control the rate at which retransmissions are sent, one disadvantage of this approach is that retransmissions occur much more slowly than they would if implemented in firmware. As a compromise, we set the retransmit count to 2 so that most retransmissions are handled by the firmware, but the driver is still informed fairly quickly when channel conditions are too poor to allow the packet to be sent.

Unfortunately, the Prism 2.5 chipset does not support per-packet transmission power control the way it does for the transmission rate. This is *not* a fundamental limitation of wireless NICs, but rather a characteristic of the Prism 2.5 firmware. We overcome this limitation as follows. Whenever we want to change the transmission power, we wait for the NIC's transmission buffers to empty. We then change the power level and queue packets that should be sent at the new power level. While this technique supports per-packet power control for packets which pass through the driver (user data), it does not allow us to set the power for 802.11 control and management packets (e.g. ACKs, RTS/CTS, beacons), handled completely in firmware; these packets will simply be sent at the power level that the card happens to be using at the time. Our approach also introduces overhead (extra

idle time) for each power level set operation. Nevertheless, it enables us to examine the basic tradeoffs resulting from changing transmission power levels.

7.2 Fixed-power rate selection algorithms

Most 802.11b implementations select transmission rate using a variation of the *Auto Rate Fallback* (ARF) algorithm [31]. ARF attempts to select the best transmission rate via in-band probing using 802.11's ACK mechanism. ARF assumes that a failed transmission indicates a transmission rate that is too high. A successful transmission is assumed to indicate that the current transmission rate is good, and that a higher rate might be possible.

Our ARF implementation works as follows. If a threshold number of consecutive packets are sent successfully, the node selects the next higher transmission rate. If a threshold number of consecutive packets are dropped, the node decrements the rate. If no traffic has been sent for a certain time, the node uses the highest possible transmission rate for the next transmission. In our implementation, the increment threshold is set to 6 successful packet transmissions, the decrement threshold to 4 dropped packets (that is, 2 notifications of transmission failure from the firmware as discussed in Section 7.1), and the idle timeout value to 10 seconds. Within the constraints of our driver-based approach, these settings are designed to approximate the Prism 2.5 firmware's implementation of ARF algorithm which appears to use an increment threshold of 6, a decrement threshold of 3, and an idle timeout of 10 seconds.

An alternative to probing the channel for the best transmission rate is to use the channel's signal-to-noise ratio (SNR) to select the optimal rate for a given SNR. While SNR-based rate selection algorithms eliminate the overhead of probing for the correct transmission rate, they face a number of practical challenges. First, card measurements of SNR can be inaccurate and may vary between different cards of the same make and model. Second, SNR measurements do not completely characterize channel degradation due to multipath interference. Finally, the information that SNR-based rate selection algorithms need is measured at the receiver, since it is the SNR at the receiver that determines whether or not a packet is received successfully. While proposals have been made for overcoming this problem by leveraging 802.11's RTS/CTS mechanism (see, for example, [11]), these solutions do not work on current hardware.

In our implementation, we overcome the last challenge by tagging each packet with channel information. Specifically, each packet contains the transmit power level used to send it, as well as the path loss and noise estimate of the last packet sent from the destination towards the sender. This allows the receiver to estimate both uplink and downlink path loss

information (both are required as asymmetry may arise due to antenna diversity).

The SNR-based algorithm that we use, *Estimated Rate Fallback* [17] (ERF), is actually a hybrid between pure SNR-based and ARF-based algorithms. It uses the path loss information to estimate the SNR with which each transmission will be received. ERF then determines the highest transmission rate that can be supported for this SNR. In addition, since SNR measurements have some uncertainty, if the estimated SNR is just below a rate selection decision boundary, ERF will try the rate immediately above the estimate best transmission rate after a given number of successful sends. Similarly, if the estimated SNR is just above a decision threshold, ERF will use the rate immediately below the estimated best transmission rate after a given number of failures. Finally, if no packets have been received from the destination for a given interval, ERF will begin to fall back towards the lowest rate until new channel information has been received. This keeps ERF from getting stuck in a state where stale channel information prevents communication, which in turn prevents obtaining new channel information.

7.3 Power and rate selection algorithms

Both the ARF and ERF algorithms use a fixed transmission power, which is typically set fairly high to maximize the chance that a node can communicate with the intended destination. However, as we discussed earlier, high power levels create significant interference which may reduce performance for other channels. Next, we discuss two algorithms that combine power and rate control to try to minimize interference. In both strategies, each transmitter attempts to reduce its power to the minimum level that allows it to reach the intended receiver at the maximum transmission rate. In essence, each sender acts socially (by reducing interference) as long as it does not cost anything (no rate reduction).

ARF can be extended naturally to support conservative power control by adding low power states above the highest rate state: At the highest rate, after a given number of successful sends, transmit power is reduced by a fixed amount. This process repeats until either the lowest transmit power is reached or the transmission failed threshold is reached. In the latter case, the power is raised a fixed amount. If failures continue, power is raised until the maximum transmit power is reached after which rate fallback begins. We call this algorithm *Power-controlled Auto Rate Fallback* (PARF).

ERF can also be easily extended to implement conservative power control as follows. First, an estimated SNR at the receiver is computed as described earlier. Now, if the estimated SNR is a certain amount (the “power mar-

gin”) above the decision threshold for the highest transmit rate, the transmit power is lowered until $estimatedSNR = decisionThreshold + powerMargin$. The *powerMargin* variable allows the aggressiveness of the power control algorithm to be tuned. We call this algorithm *Power-controlled Estimated Rate Fallback* (PERF).

8 Performance evaluation

We now discuss the results of some basic experiments we conducted to verify that current commodity hardware is able to achieve performance improvements when using transmit power control. These experiments use the modified HostAP driver and Prism 2.5 802.11b cards described previously.

8.1 Interference test

We first measured the interference to an aggressive TCP flow from a bandwidth-limited TCP flow for the ARF, ERF, and PERF algorithms discussed previously. Note that we did not use the PARF algorithm. In our initial experimentation with PARF, we found that its behavior was quite unstable. This may be because power decrease decisions by the “receiver” could result in the ACK-half of a data transmission/ACK exchange failing due to insufficient power. This could, in turn, result in the sender increasing its power even if it was the fault of the receiver’s power level setting. While this could also happen to PERF (since it also considers packet losses), PERF largely avoids this overreaction to packet losses by basing most of its power and rate selection on the SNR of the interactions. This makes PERF react more slowly to transmission failures. As a result, PERF always behaved more stably and performed significantly better than PARF and we did not consider PARF any further in our evaluation.

Our experimental topology is shown in Fig. 14. Two AP-client pairs communicate in a laboratory setting where all nodes communicate via coaxial cables. Attenuators are placed on the cables to control the attenuation between nodes. The “victim” pair repeatedly executes a TCP throughput benchmark from the AP to the client and always runs the ERF algorithm, but with power set to 0 dBm instead of the default 23 dBm. The “aggressor” pair executes a rate-limited transfer of 1.2 Mbps from the AP to the client. The power




Fig. 14 Laboratory interference test topology




Fig. 15 Lab interference test

and rate selection algorithm used by the aggressor pair was varied to measure the ability of the algorithms to reduce the interference experienced by the victim pair.

As shown in Fig. 15(a), the power reduction used by PERF nearly completely eliminated the interference experienced by the victim pair. (The error bars show 95% confidence intervals.)

Figure 15(b) shows the same test but with the aggressor transfer running at unlimited speed and with 106 dB of loss between the pairs. This represents an extreme situation that may occur to some degree even in current fixed power networks due to heterogeneous transmit power levels. Clearly, in this case, the victim’s communication under ARF and ERF is practically zero. The poor performance for ARF and ERF is a result of asymmetric carrier sense [24]. In this experiment, the victim pair is able to sense transmissions of the aggressor pair and defer transmissions to prevent collisions. However, the aggressor pair is not able hear the lower powered transmissions of the victim. Therefore, it transmits packets as quickly as possible, essentially interfering with any transmission between the victim pair. Note that the aggressor transmissions are nearly always received successfully — the aggressor pair

always obtains a throughput of approximately 3.5 Mbps in each setting. PERF alleviates this situation by reducing the power of the aggressor since the aggressor AP-client link is overprovisioned in terms of power. In this experiment, PERF is able to nearly completely isolate the pairs from each other.

Note that when using power control, care must be taken so that this asymmetric carrier sense situation is not actually introduced inadvertently. That is, an AP-client link should not reduce power so much that it becomes overwhelmed by neighboring uncooperative high power nodes. This is undesirable since we assume that APs and clients do not desire to sacrifice (much) performance when lowering transmit power.

To show that a similar situation can occur outside of the laboratory, we then repeated the rate-limited interference test in a residential setting where the nodes communicated over the air as shown in Fig. 16(a) instead of over coaxial cables as they did in the laboratory tests. Figure 16(b) shows the results of this test. ARF and ERF experienced similar performance degradation due to interference. PERF was able to reduce this degradation by about 50%.


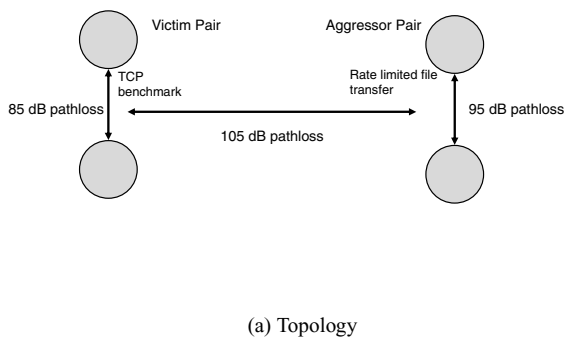


Fig. 16 Home interference test

9 Discussion

In our evaluation, we considered protocols that reduce power as long as transmission rate was unaffected. One possibility to reduce interference even further is to consider algorithms that allow the transmission rate to be reduced. In one such strategy that we also considered in our implementation, called *Load-sensitive, Power-controlled Estimated Rate Fallback (LPERF)*, transmitters reduce their power even if it reduces their transmission rate. Specifically, they reduce their power as long as the resulting transmission rate is sufficient to support their actual traffic demands.

While potentially increasing network throughput, LPERF involves a subtle tradeoff between the scope of the interference (controlled by power level) and the duration of interference (determined by the transmission rate). Our analysis in Section 5 indicates that, at least for simple topologies, achieving the highest possible AP density requires using an LPERF-like approach where lower transmit rates are used to enable lower transmit power settings. In practice, we found that achieving good performance and interference reduction using the LPERF technique can be challenging.

First, we must be sure that actual demand does not require a higher transmission rate than the one we select. This can be difficult since demand may be a function of the transmission rate used. In addition, since the wireless medium is shared, the ability to satisfy traffic demand at a particular transmission rate depends on the particular fraction of the shared medium that a node receives. Assuming a fair distributed MAC protocol, the size of this fair share depends on the demands of the other nearby nodes. An additional complication is that the set of nearby nodes changes with the transmission power used since the power determines the range of any transmission. In order to address this need, LPERF incorporates techniques to continuously monitor link traffic demand and medium utilization. Tuning such techniques to adapt quickly to changes in demand is an open research question.

Second, we must have accurate measurements of received signal strength, noise, and transmit power from all nodes in the area. For the hardware we used in our implementation, 4 dB of variance in RSS and noise estimates is typical. As the entire range of SNR thresholds for 802.11b transmission rates is only 9 dB, this can be a significant issue.

In summary, the design of algorithms that consider lowering transmission rates is a challenging task. In the future, we expect that the range of transmit rates supported by wireless cards will greatly increase. This greatly helps such algorithms by both creating more situations where the maximum link bandwidth is not used and by providing a much wider range of power levels that such algorithms can employ. In addition, we also have anecdotal evidence that vendors are providing more accurate RSS and noise estimators in newer

cards. As a result, we believe that LPERF-like algorithms are a promising direction.

10 Summary

A chaotic network consists of a set of co-located wireless nodes owned and controlled by different people or organizations. Its main characteristic is that the deployment is largely unplanned and unmanaged. In this paper, we studied important characteristics of chaotic wireless networks. We used measurements from several cities to show that chaotic networks can be quite dense and we used trace-driven simulations to show that the performance of end-clients can suffer significantly in these dense chaotic wireless deployments. We also presented and evaluated automated power control and rate adaptation algorithms that reduce interference among neighboring nodes while ensuring robust endclient performance. Specifically, we showed how the LPERF algorithm, that reduces transmission power as much as possible without reducing transmission rate, can improve aggregate throughput significantly in our small testbed. We also discussed how further improvements may be possible in denser deployments.

References

1. D. Aguayo, J. Bicket, S. Biswas, G. Judd and R. Morris, Link-level measurements from an 802.11b Mesh network, in: *Proceedings ACM SIGCOMM* (2004) pp. 121–132.
2. A. Akella, R. Karp, S. Seshan, S. Shenker and C. Papadimitriou, Selfish behavior and stability of the internet: A game-theoretic analysis of TCP, in: *Proceedings ACM SIGCOMM* (2002) pp. 117–130.
3. Alcatel, AirView Software, <http://www.alcatel.com>.
4. Y. Cheng, Y. Chawathe, A. LaMarca and J. Krumm, Accuracy characterization for metropolitan-scale Wi-Fi localization, in: *Proceedings MobiSys* (2005) pp. 233–245.
5. D. Chiu and R. Jain, Analysis of the increase/decrease algorithms for congestion avoidance in computer networks, *Computer Networks and ISDN Systems* 17 (1989) 1–14.
6. D. Clark, C. Partridge, J. Ramming and J. Wroclawski, A knowledge plane for the internet, in: *Proceedings ACM SIGCOMM* (2003) pp. 3–10.
7. Datacomm Research, New Datacomm Research Report: Wireless LAN Equipment Shipments to Triple Within Five Years, <http://www.tmcnet.com/usubmit/2005/Feb/1120138.htm>.
8. R. Draves, J. Padhye and B. Zill, Comparison of routing metrics for static multi-hop wireless networks, in: *Proceedings ACM SIGCOMM* (2004) pp. 133–144.
9. R. Droms, Dynamic host configuration protocol, Technical report, Internet Engineering Task Force. RFC 2131.
10. A. Hills, Large-scale wireless LAN design, *IEEE Communications* 39 (2001) 98–104.
11. G. Holland, N. Vaidya and P. Bahl, A rate-adaptive MAC protocol for multi-hop wireless networks, in: *Proceedings ACM MobiCom* (2003) pp. 236–251.

12. IEEE, OUI and Companyid assignments, <http://standards.ieee.org/regauth/oui/oui.txt>.
13. InStat/MDR, 3Q 2004 WLAN Market Share Report, <http://www.instat.com/tr/nrep/2004/IN0401429WL.htm>.
14. Intego, WI-Fi Locator, <http://www.intego.com/wiFiLocator/>.
15. Intel Research Seattle, Place Lab: A Privacy-Observant Location System, <http://placelab.org/>.
16. JIwire, WI-FI Hotspot locator, <http://jiwire.com>.
17. G. Judd and P. Steenkiste, Using emulation to understand and improve wireless networks and applications, in: *Proceedings USENIX NSDI (2005)*.
18. V. Kawadia and P.R. Kumar, Principles and protocols for power control in Ad Hoc networks, *IEEE Journal on Selected Areas in Communications* 1 (2005) 5–16.
19. R. Kompella and A. Snoeren, SPARTA: Scheduled power and rate adaptation, in: *ACM SenSys (2003)*.
20. B. Mah, An empirical model of HTTP network traffic, in: *Proceedings IEEE Infocom (1997)* pp. 592–600.
21. J. Malinen, Host AP Driver, <http://hostap.epitest.fi/>.
22. Propagate Networks, Autocell, <http://www.propagatenetworks.com/product/>.
23. D. Qiao, S. Choi, A. Jain and K. Shin, MiSer: An optimal low-energy transmission strategy for IEEE 802.11a/h, in: *Proceedings ACM MobiCom (2003)* pp. 161–175.
24. A. Rao and I. Stoica, An overlay MAC layer for 802.11 networks, in: *Proceedings MobiSys (2005)* pp. 135–148.
25. T. Rappaport, *Wireless Communications: Principles and Practice* (Prentice-Hall, Englewood Cliffs, NJ, 1996).
26. J. Rexford, A. Greenberg, G. Hjalmtysson, A. Myers, G. Xie, J. Zhan and H. Zhang, Network-wide decision making: Toward a wafer-thin control plane, in: *Proceedings ACM SIGCOMM HotNets-III*.
27. B. Sadeghi, V. Kanodia, A. Sabharwal and E. Knightly, Opportunistic media access for multi-rate Ad Hoc networks, in: *Proceedings ACM MobiCom (2002)* pp. 24–35.
28. A. Santhanam and R. Cruz, Optimal routing, link scheduling and power control in multi-hop wireless networks, in: *Proceedings IEEE Infocom (2003)* pp. 702–711.
29. Strix Systems, AccessOne/Network OWS, http://www.strixsystems.com/products/products_main.asp.
30. S. Thomson and T. Narten, IPv6 Stateless Address Autoconfiguration, Technical report, Internet Engineering Task Force. RFC 2462.
31. V. van der Vegt, Auto Rate Fallback, <http://www.phys.uu.nl/vdvegt/docs/gron/node24.html>.
32. Wi-Fi-Zones.com, Find more hotspot locations, <http://www.wi-fi-zones.com>.
33. Wifi Maps.com, Wardriving Maps and Hotspot Locator, <http://www.wifi-maps.com>.
34. X. Zeng, R. Bagrodia and M. Gerla, Global Mobile Information Systems Simulation Library, <http://pcl.cs.ucla.edu/projects/glomosim/>.
35. ZeroConf, IETF Zero configuration Networking Working Group, <http://www.ietf.org/html.charters/zeroconf-charter.html>.



Aditya Akella obtained his Ph.D. in Computer Science from Carnegie Mellon University in September 2005. He obtained a B.Tech in Com-

puter Science and Engineering from IIT Madras in May 2000. Currently, Dr. Akella is a post-doctoral associate at Stanford University. He will join the Computer Sciences faculty at the University of Wisconsin-Madison in Fall 2006. Dr. Akella's research interests include Internet Routing, Network Protocol Design, Internet Security, and Wireless Networking. His web page is at <http://www.cs.wisc.edu/~akella>.



Glenn Judd is a Computer Science Ph.D. candidate at Carnegie Mellon University. His research interests include wireless networking and pervasive computing. He has an M.S. and B.S. in Computer Science from Brigham Young University.



Srinivasan Seshan is currently an Associate Professor and holds the Finmeccanica chair at Carnegie Mellon University's Computer Science Department. Dr. Seshan received his Ph.D. in 1995 from the Computer Science Department at University of California, Berkeley. From 1995 to 2000, Dr. Seshan was a research staff member at IBM's T.J. Watson Research Center. Dr. Seshan's primary interests are in the broad areas of network protocols and distributed network applications. In the past, he has worked on topics such as transport/routing protocols for wireless networks, fast protocol stack implementations, RAID system design, performance prediction for Internet transfers, Web server benchmarking, new approaches to congestion control, firewall design and improvements to the TCP protocol. His current work explores new approaches in overlay networking, sensor networking, online multiplayer games and wide-area Internet routing. His web page is at <http://www.cs.cmu.edu/~srini>.



Peter Steenkiste is a Professor of Computer Science and of Electrical and Computer Engineering at Carnegie Mellon University. His research interests include networking, distributed systems, and pervasive computing. He received an M.S. and Ph.D. in Electrical Engineering from Stanford University and an Engineering degree from the University of Gent, Belgium. You can learn more about his research from his home page <http://www.cs.cmu.edu/~prs>.