Self-Management in Chaotic Wireless Deployments

Aditya Akella Glenn Judd Srinivasan Seshan Peter Steenkiste

Carnegie Mellon University

{aditya, glennj, srini+, prs}@cs.cmu.edu

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.

Categories and Subject Descriptors

C.2 [**Computer Systems Organization**]:
Computer-Communication Networks;
C.2.1 [**Computer-Communication Networks**]:
Network Architecture and Design;
Wireless communication

General Terms

Measurement, Performance, Experimentation

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.

Keywords

access points, interference, power control, channel assignment

1. INTRODUCTION

Wireless data 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 a result, 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 resources that are 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 [18]. However, the rapid deployment of cheap 802.11 hardware and other personal wireless technology (2.4GHz 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 [21] and that the sales of WiFi equipment will triple by 2009 [14]. 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 each setting up one or a small number of APs. This type of spontaneous deployment results in highly variable densities of wireless nodes and APs and, in some cases, these densities can become very high (e.g. urban environments, apartment buildings). Moreover, 802.11 nodes have to share the spectrum with other networking technologies (e.g., Bluetooth, UWB) 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 [22] 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 levels 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 (11Mbps 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. We believe that 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

In this section, 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 WifiMaps [8], Wi-Fi-Zones.com [7] and JIWire.com [6]. Several vendors also market products targeted at locating wireless networks while on the go (see for example, Intego WiFi Locator [5]. Among research studies, the Intel Place Lab project [22] [11] maintains a database of up to 30,000 802.11b APs from several US cities. In this paper, we use hot-spot data from WifiMaps.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 [36, 34] and have been widely deployed [16], a number of interesting research problems in simplifying network management still remain (e.g., [13, 31]). 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 [3], Strix Systems' Access/One Network [1] and Alcatel OmniAccess' AirView Software [2]. 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 load management and effective 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 OmniAcess 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, in the past, several rate adaptation mechanisms that leverage the multiple rates supported by 802.11 have been proposed. For example, Sadeghi et al. [32] study new multi-rate adaptation algorithms to improve throughput perfor-

Data set	Collected on	No. of APs	Stats collected per AP
Place Lab	Jun 2004	28475	MAC, ESSID, GPS coordinates
WifiMaps	Aug 2004	302934	MAC, ESSID, Channel
Pittsburgh Wardrive	Jul 2004	667	MAC, ESSID, Channel supported rates, GPS coordinates

Table 1: Characteristics of the data sets

mance in ad hoc networks. Our rate control algorithms, in contrast, are designed specifically to work well in conjunction with power control. However, it is possible to extend past algorithms such as [32] 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, for example, [28, 25]). 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. [24, 15, 33, 20]. 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. In this section, we present preliminary observations of the density of APs in these metropolitan areas, as well as typical usage characteristics, such as the channels used for transmission and common vendor types.

3.1 Measurement Data Sets

We use three 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 [22] [11] in June 2004. The Place Lab software allows commodity hardware clients like notebooks, PDAs and cell phones to locate themselves by listening for radio beacons such as 802.11 APs, GSM cell phone towers, and fixed Bluetooth devices.
2. **WifiMaps**: The WifiMaps.com website [8] provides a 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.
3. **Pittsburgh Wardrive**: This data set was collected on July 29, 2004, as part of a small-scale wardriving effort which covered 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.

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 802.11g 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 50m, 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 Boston and San Diego, for example, a particular wireless AP suffers interference from about 80 other APs deployed in close proximity.

In Figure 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

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

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

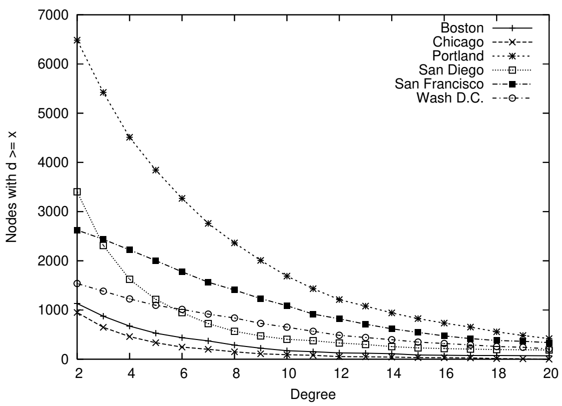

Figure 1: Distribution of AP degrees (Place Lab data set)

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 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.

3.2.2 802.11 Usage: Channels

Table 3 presents the distribution of channels used by APs in the WiFiMaps data set. 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 14% 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.

3.2.3 802.11b vs. 802.11g

The Pittsburgh wardrive 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

Channel	Percentage of APs
1	3.04
2	12.29
3	3.61
4	1.03
5	1.13
6	41.15
7	1.75
8	1.12
9	1.31
10	3.42
11	11.04

Table 3: Channels employed by APs in the Wifimaps data set.

classified APs, or about 93, are 802.11g. Given the relatively recent standardization of 802.11g (June 2003), these measurements suggest that new wireless technology gets deployed relatively quickly.

3.2.4 Vendors and AP Management Support

Vendor	Percentage of APs
Total classified	98
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

Table 4: Popular AP vendors (Wifimaps data set)

To determine popular AP brands, we look up the MAC addresses available in the Wifimaps data set against the IEEE Company_id assignments [4] to classify each AP according to the vendor. For the APs that could be classified in this manner (2% of the APs in the Wifimaps data set did not have a matching vendor name), the distribution of the vendors is shown in Table 4. 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 4. 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.

4. IMPACT ON END-USER PERFORMANCE

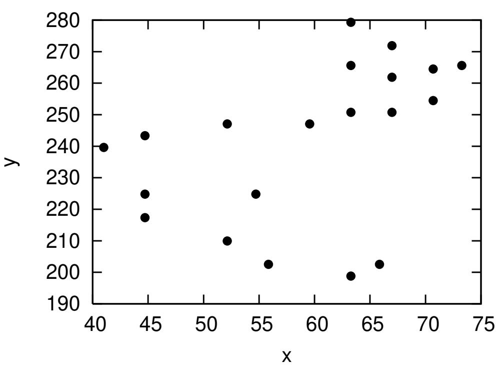

(a) 20-node topology

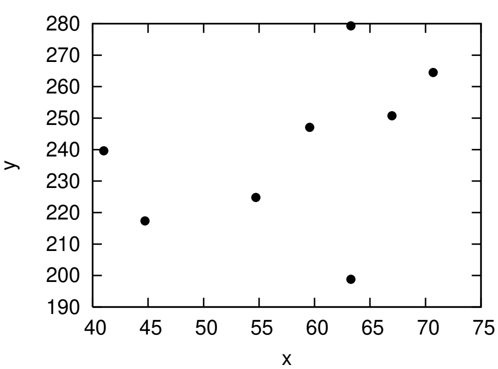

(b) 8-node sub-topology

Figure 2: Simulation topologies derived from the Metropolis wardrive data set; the units on the x and y axis are in meters. Figure (a) shows the 20-node topology derived from the Metropolis Wardrive data set. Figure (b) shows a sub-topology of 8 nodes that were all assigned the same channel by a static optimal channel assignment algorithm when applied to the 20-node topology in (a).

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 [17]. We simulated the deployment topology shown in Figure 2(a), obtained during a portion of the Pittsburgh wardrive. There are 20 APs in this topology. 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 3.
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 15dBm, unless otherwise specified (This is the default setting in most commercial APs).
6. All APs transmit at a single rate, 2Mbps (there is no multi rate support in GloMoSim). At these settings, the transmission and interference ranges are 31m and 65m, respectively.
7. RTS/CTS is turned off. This is the default setting in most commercial APs.
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 [30].

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 [26]. 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 5s and 20s (We also simulated HTTP workloads with 10s, 30s and 60s sleep times. The results are qualitatively similar and are omitted for brevity). The average load offered by the HTTP client is 83.3Kbps for a 5s sleep time, and 24.5Kbps for a 20s sleep time. There is no other interfering traffic in the `http` workload.

The second set of workloads, called $\texttt{comb-ftp}_i$, 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 0.89Mbps. We simulate either set of workloads for 300s.

4.1 Interference at Low Client-Densities and Traffic Volumes

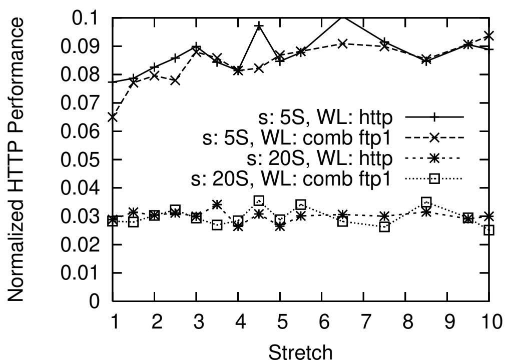

(a) HTTP, $D = 1$

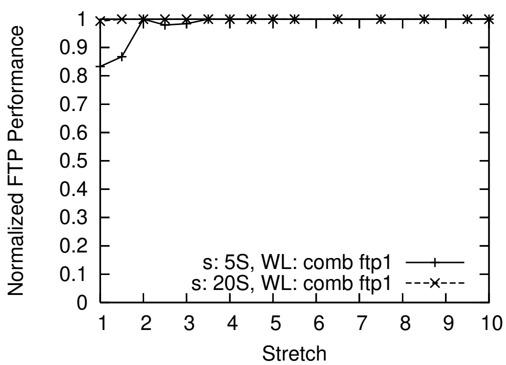

(b) FTP, $D = 1$

Figure 3: 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).

First, we conduct simulations with the `http` and `comb-ftp`$_1$ workloads, and low client densities ($D = 1$). The results are shown in Figure 3. 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 a factor of l larger than 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 topology in Figure 2(a), 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 (Figure (a)) or FTP flows (Figure (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., 0.89Mbps. 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 5s), 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 20s), 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 Figure 3(b). When the HTTP component of this workload is aggressive ($s = 5s$), 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 the traffic volume of the clients—to create more aggressive interference settings.

4.2 Impact of Client Densities and Traffic Load

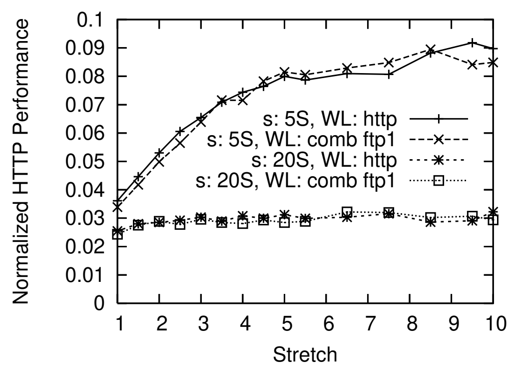

(a) HTTP performance

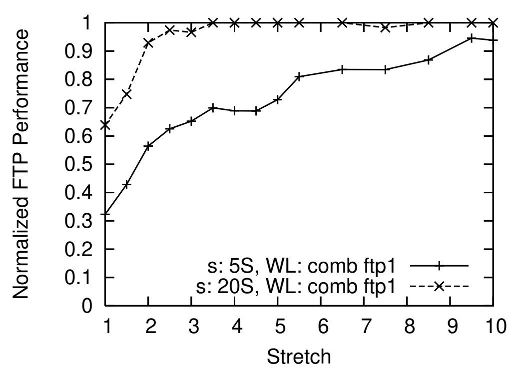

(b) FTP performance

Figure 4: Average performance of HTTP and FTP flows at greater client densities ($D = 3$).

Impact of client density. Figures 4(a) and (b) show the average performance of the individual HTTP and FTP sessions, respectively, in the `comb-ftp`$_1$ and `http` workloads, for a high number of clients associated per AP ($D = 3$). The performance of both HTTP and FTP flows suffers significantly under high client densities: From Figure 4(a), HTTP performance is lowered by about 65% (compare $stretch = 1$ with $stretch = 10$) due to interference between aggressive HTTP flows ($s = 5s$). The same is true for the performance of the FTP flow in Figure 4(b). For a less aggressive HTTP component ($s = 20s$) the performance of the HTTP flows is 20% inferior, while the FTP flow suffers by about 36%.

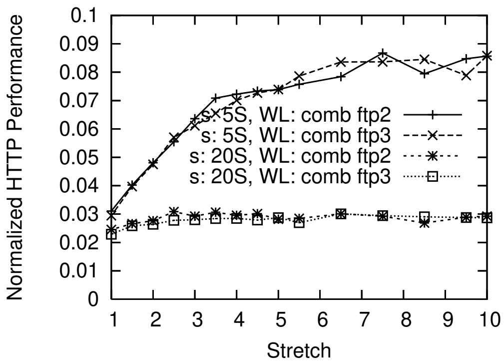

(a) HTTP performance

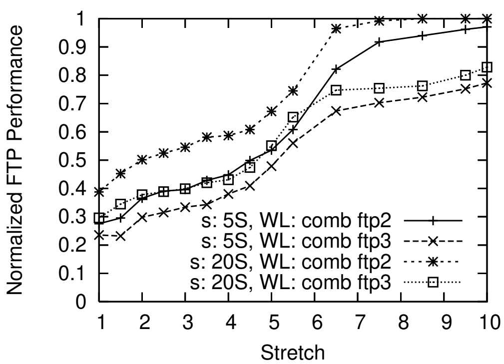

(b) FTP performance

Figure 5: Average performance of HTTP and FTP flows at higher competing FTP traffic levels (i.e., the comb-ftp{2, 3} workloads) and for $D = 3$.

Impact of traffic volume. Figures 5(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 $\texttt{comb-ftp}_2$ and $\texttt{comb-ftp}_3$ 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 $s = 20s$ in Figure 5(b)).

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 3. 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 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 20-node topology of Figure 2(a), 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. As an illustration, Figure 2(b) shows the lay-out of APs that were all assigned channel 1 by this scheme.

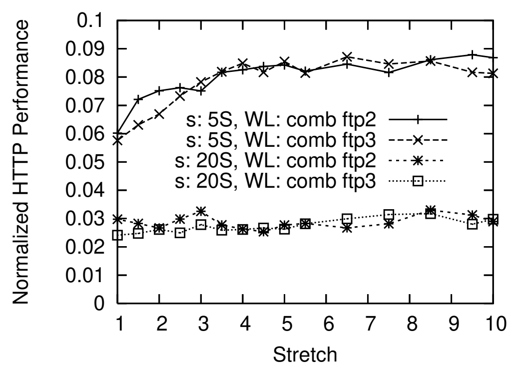

(a) HTTP performance

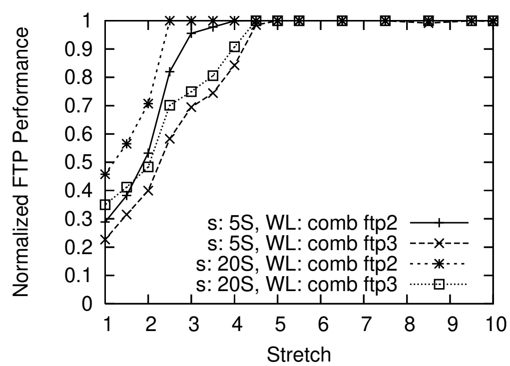

(b) FTP performance

Figure 6: Performance of HTTP and FTP flows with optimal static assignment of APs to the three non-overlapping channels. The transmit power level is set at 15dBm, corresponding to a reception range of 31m.

The performance of HTTP and FTP flows in these simulations are shown in Figure 6(a) and (b), respectively. The average performance of both HTTP and FTP flows improves significantly. Comparing with Figures 5(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 8 APs in Figure 2(b) to use a power level of 3dBm. This yields a transmission range of 15m, which is half the range from using the default power level of 15dBm. Next, we show how this improves HTTP and FTP performance, as well as the total network capacity.

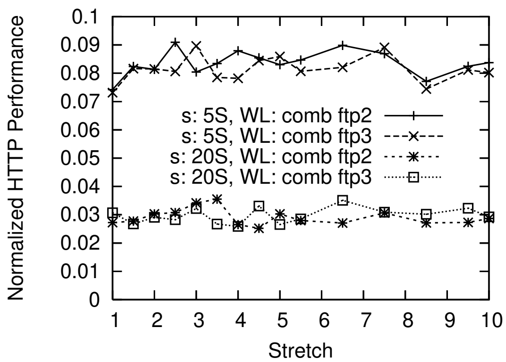

(a) HTTP performance (range = 15m)

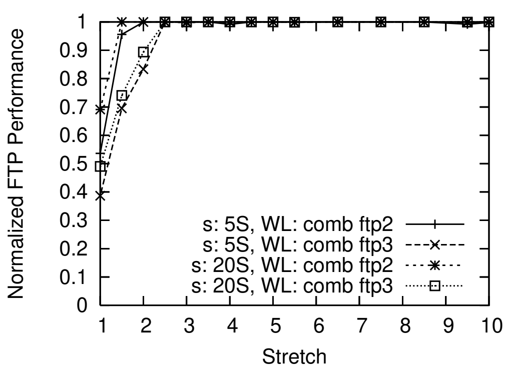

(b) FTP performance (range = 15m)

Figure 7: Performance of HTTP and FTP flows with optimal static channel assignment of APs. The transmit power level is set at 3dBm, corresponding to a reception range of 15m.

Improvement in application performance. The performance results for HTTP and FTP flows in these simulations are shown in Figures 7(a) and (b) respectively. Compared with Figures 6(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 Figure 8 we show how transmit power control improves user performance for

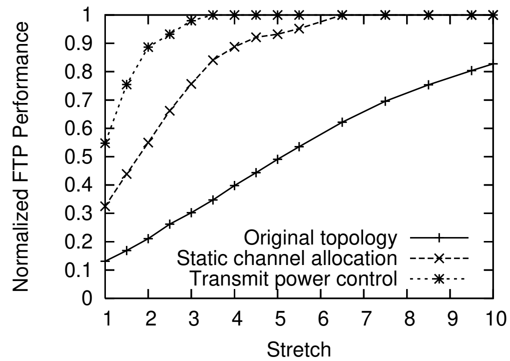

(a) FTP performance

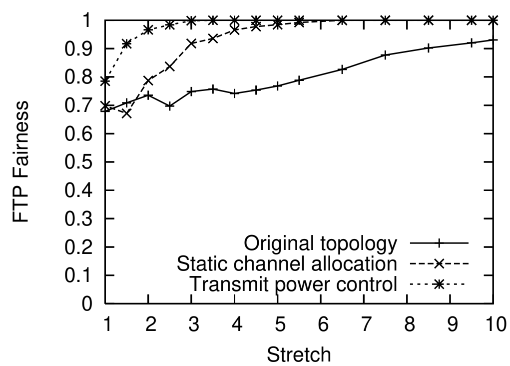

(b) Fairness

Figure 8: 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.

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 Figure 8(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 Figure 8(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 [12] 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 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 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

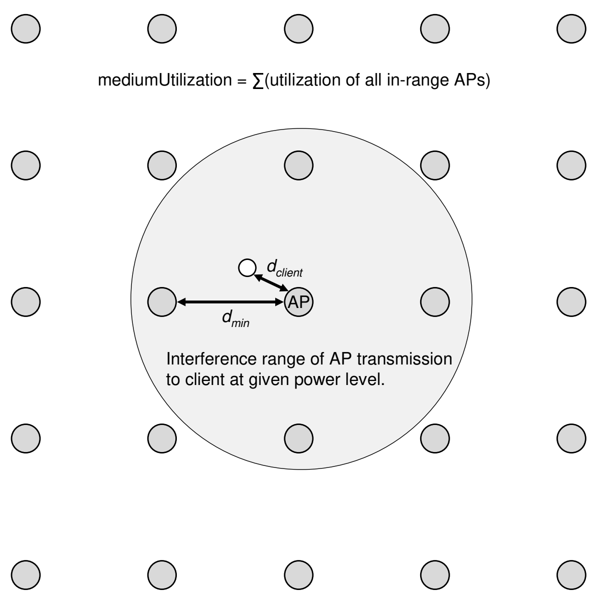

Figure 9: Computing minimum AP spacing for a grid topology

In this section we use a simple model of wireless communication applied to a two-dimensional grid topology as shown in Figure 9 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 [30] 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 5 and 6 (based on measurements presented in [9]) we then determine the maximum transmission rate that can be used for the AP-client link and the corresponding maximum throughput for that rate.

Rate (Mbps)	Minimum SNR (dB)
1	3
2	4
5.5	8
11	12

Table 5: Minimum required SNR for Prism 2.5

Rate (Mbps)	Throughput (Mbps)
1	0.85
2	1.7
5.5	3.5
11	4.9

Table 6: Maximum 802.11b throughput

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 10 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 -15dBm power on the x-axis; using the 1.1 Mbps load line, this translates into a 20m AP distance on the y-axis. This implies that if you want to transmit at 1.1Mbps from each AP at -15dBm, the APs must be at least 20m 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 5.5Mbps to communicate with its client. For the simulations in the 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.

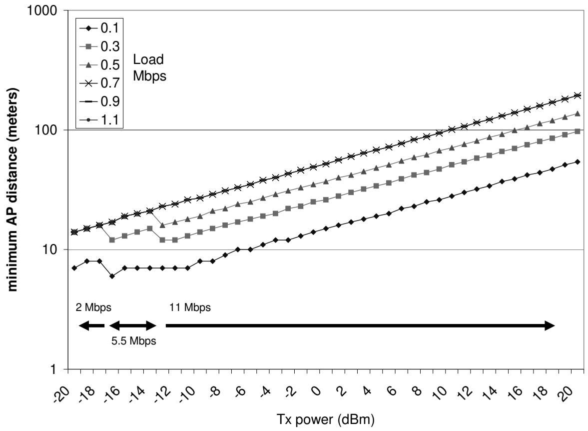

Figure 10: Minimum AP distance vs. Tx power ($d_{client} = 10m$)

We draw a few key conclusions from Figure 10. 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 [10]. 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 a AP load-balancing which arise in such environments. We leave the extension of our design to campus deployments 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 [27] 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 transmission 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 still 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 first wait for the NIC's transmission buffers to empty. We then change the power level and queue one or more 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 transmit power for 802.11 control and management packets (e.g. ACKs, RTS/CTS, beacons) which are 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 [35]. 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 transmission 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 transmission 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, [19]), 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* [23] (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. The power level 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. In this section, 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. That is, 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 transmit power is raised a fixed amount. If failures continue, the transmit 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 margin") 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 it's 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.

The topology of our experiment is shown in Figure 11. 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

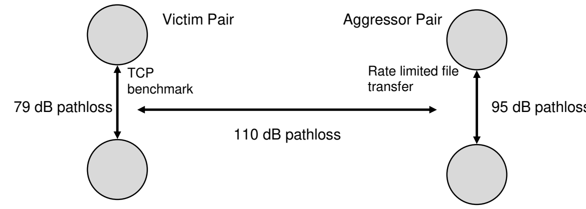

Figure 11: Laboratory Interference Test Topology

between each node. The "victim" pair repeatedly executes a TCP throughput benchmark from the AP to the client and always runs the ERF algorithm, but with power manually set to 0 dBm instead of the default power of 23 dBm. The "aggressor" pair executes a rate-limited transfer of 1.2 Mbps from the AP to the client. The power 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.

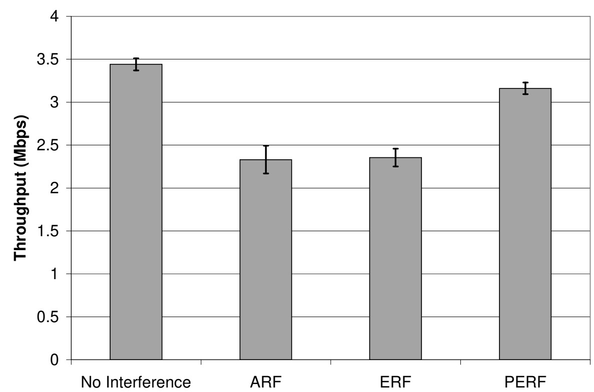

Figure 12: Lab Interference Test

As shown in Figure 12, the power reduction used by PERF nearly completely eliminated the interference experienced by the victim pair. (The error bars show 95% confidence intervals.)

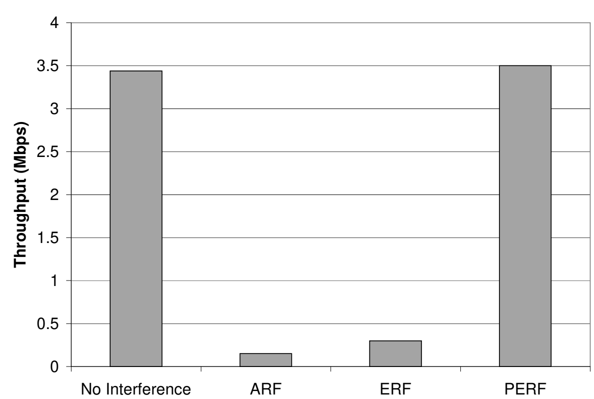

Figure 13: Lab Interference Test - Unlimited Rate

Figure 13 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 [29]. 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.

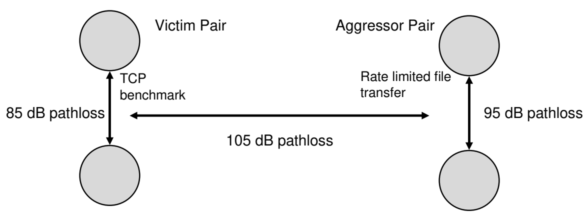

Figure 14: Home Interference Test Topology

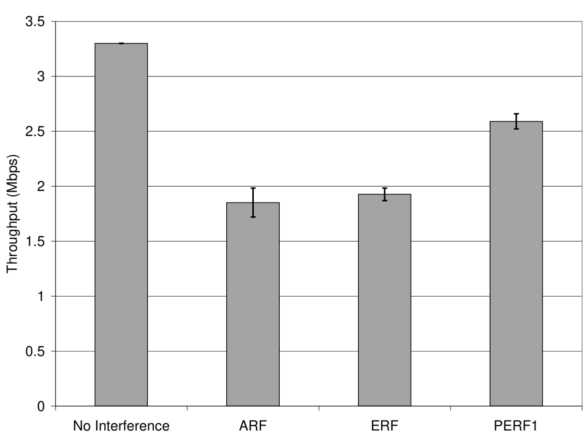

Figure 15: Home Interference Test

To demonstrate 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 Figure 14 instead of over coaxial cables as they did in the laboratory tests. Figure 15 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%.

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 total network throughput, LPERF involves a fairly 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 like the one analyzed, achieving the highest possible access point 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 end-client performance. Specifically, we showed how the PERF 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.

Acknowledgment

We would like to thank Drew Celley and Eric Blevins of WifiMaps.Com for providing us access to their vast database. We are grateful to Yatin Chawathe for pointing us to the Intel Place Lab database. We thank UCLA's Parallel Computing Lab for GloMoSim. We also thank Rahul Dhar for help with our live experiments. Feedback from Suman Banarjee, Carl Gunter, Brad Karp and David Wetherall helped improve this paper substantially. Finally, we thank our shepherd, Victor Bahl, and our anonymous reviewers for their valuable feedback and suggestions.

11. REFERENCES

[1] AccessOne/Network OWS. http://www.strixsystems.com/products/products_main.asp.

[2] Alcatel AirView Software. http://www.alcatel.com.

[3] Autocell. http://www.propagatenetworks.com/product/.

[4] IEEE OUI and Companyid assignments. http://standards.ieee.org/regauth/oui/oui.txt.

[5] Intego WI-Fi Locator. http://www.intego.com/wiFiLocator/.

[6] WI-FI Hotspot locator. http://jiwire.com.

[7] Wi-Fi-Zones.com - Find more hotspot locations. http://www.wi-fi-zones.com.

[8] WiFiMaps.com - Wardriving Maps and Hotspot Locator. http://www.wifimaps.com.

[9] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *Proceedings of the ACM SIGCOMM Conference on Network Architectures and Protocols*, Portland, August 2004.

[10] 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 of the SIGCOMM '02 Symposium on Communications Architectures and Protocols*, Pittsburgh, PA, August 2002. http://www.cs.cmu.edu/~aditya/papers/cc-gametheory.pdf.

[11] Y. Cheng, Y. Chawathe, A. LaMarca, and J. Krumm. Accuracy characterization for metropolitan-scale Wi-Fi localization. In *Proceedings of MobiSys05*, Seattle, WA, June 2005.

[12] D. Chiu and R. Jain. Analysis of the Increase/Decrease Algorithms for Congestion Avoidance in Computer Networks. *Computer Networks and ISDN Systems*, 17(1):1–14, June 1989.

[13] D. Clark, C. Partridge, J. C. Ramming, and J. Wroclawski. A Knowledge Plane for the Internet. In *Proceedings of ACM SIGCOMM 2003*, Karlsruhe, Germany, August 2003.

[14] Datacomm Research. New Datacomm Research Report: Wireless LAN Equipment Shipments to Triple Within Five Years. http://www.tmcnet.com/usubmit/2005/Feb/1120138.htm, 2005.

[15] R. Draves, J. Padhye, and B. Zill. Comparison of Routing Metrics for Static Multi-Hop Wireless Networks. In *SIGCOMM'04*, Portland, August 2004. ACM.

[16] R. Droms. Dynamic Host Configuration Protocol. Technical report, Internet Engineering Task Force, March 1997. RFC 2131.

[17] Global Mobile Information Systems Simulation Library. http://pcl.cs.ucla.edu/projects/glomosim/.

[18] A. Hills. Large-Scale Wireless LAN Design. *IEEE Communications*, 39(11):98–104, November 2001.

[19] G. Holland, N. Vaidya, and P. Bahl. A Rate-Adaptive MAC Protocol for Multi-hop Wireless Networks. In *Proceedings of MobiCom2001. Rome, Italy*, September 2001.

[20] G. Holland, N. Vaidya, and P. Bahl. A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks. In *Proceedings of MobiCom 2003*, Rome, July 2003. ACM.

[21] Instat/MDR. 3Q 2004 WLAN Market Share Report. http://www.instat.com/r/nrep/2004/IN0401429WL.htm.

[22] Intel Research Seattle. Place Lab: A Privacy-Observant Location System. http://placelab.org/, 2004.

[23] G. Judd and P. Steenkiste. Using Emulation to Understand and Improve Wireless Networks and Applications. In *Proceedings of NSDI 2005*, Boston, MA, May 2005.

[24] V. Kawadia and P. R. Kumar. Principles and Protocols for Power Control in Ad Hoc Networks. *IEEE Journal on Selected Areas in Communications*, 2005.

[25] R. R. Kompella and A. C. Snoeren. SPARTA: Scheduled Power and Rate Adaptation. In *ACM SenSys 2003*.

[26] B. A. Mah. An Empirical Model of HTTP Network Traffic. In *Proceedings of IEEE INFOCOM 1997*, April 1997.

[27] J. Malinen. Host AP Driver. http://hostap.epitest.fi/.

[28] D. Qiao, S. Choi, A. Jain, and K. Shin. MiSer: An Optimal Low-Energy Transmission Strategy for IEEE 802.11a/h. In *ACM MobiCom 2003*.

[29] A. Rao and I. Stoica. An overlay MAC layer for 802.11 networks. In *Proceedings of MobiSys05*, Seattle, WA, June 2005.

[30] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice-Hall, Englewood Cliffs, NJ.

[31] J. Rexford, A. Greenberg, G. Hjalmtysson, D. M. A. Myers, G. Xie, J. Zhan, and H. Zhang. Network-wide Decision Making: Toward A Wafer-Thin Control Plane. In *HotNets-III*, San Diego, CA, November 2004.

[32] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly. Opportunistic Media Access for Multi-Rate Ad Hoc Networks. In *Proceedings of MobiCom 2002*, Atlanta, GA, October 2002. ACM.

[33] A. Santhanam and R. Cruz. Optimal Routing, Link Scheduling and Power Control in Multi-hop Wireless Networks. In *Infocom '03*. IEEE, March 2003.

[34] S. Thomson and T. Narten. IPv6 Stateless Address Autoconfiguration. Technical report, Internet Engineering Task Force, December 1998. RFC 2462.

[35] V. van der Vegt. Auto Rate Fallback. `http://www.phys.uu.nl/~vdvegt/docs/gron/node24.html`.

[36] IETF Zero Configuration Networking (zeroconf) Working Group. `http://www.ietf.org/html.charters/zeroconf-charter.html`, 2000.