Wi-Fi Networks are Underutilized

Ramya Raghavendra[†], *Jitendra Padhye*[‡], *Ratul Mahajan*[‡] *and Elizabeth Belding*[†]

[†]*U. of California, Santa Barbara*, [‡]*Microsoft Research*

1. INTRODUCTION

We recently learned that Microsoft's IT department was hesitating to upgrade its Wi-Fi infrastructure to the new, 802.11n-compliant equipment. 802.11n is slated to have 2-4 times the capacity of the currently prevalent 802.11a/g standard. The source of this hesitation was their observation that the existing 802.11 a/g network was significantly underutilized, implying that the value of the upgrade would be minimal.

We were intrigued by this observation. Some of our recent research [30, 3] has been (partially) motivated by the thesis that Wi-Fi networks are growing ever-more popular, and would soon face a capacity crunch. In fact, much of recent research work on Wi-Fi networks [26, 18, 30] has been motivated by this vision.

However, these papers, including ours, offer little justification for espousing this belief. We could not find any work that had systematically studied utilization of Wi-Fi networks and made the case for additional capacity, either in the short- or the long-term.

This is not to say that performance of deployed wireless networks has not been well-studied. Various characteristics of wireless networks such as coverage [15, 10], loss rate [32], usage patterns [36, 36, 27], user mobility [22, 19] etc. have been studied by many researchers in diverse environments such as home networks [32], public hot-spots [17], corporate deployments [15, 5] and large meetings [24, 23, 34].

However, none of these studies specifically focus on medium utilization, which is the metric for determining whether the network is over or under subscribed. The medium utilization is affected not just by the traffic carried by the network, but also by interference from other, nearby Wi-Fi networks, as well as electronic equipment such as microwave ovens.

Thus, we decided to conduct a study of utilization of Wi-Fi networks in a broad range of environments. We study corporate offices, coffee shops, university buildings, houses, as well as densely occupied apartment complexes and student dormitories. In addition to our own measurements we also analyze the data available in the CRAWDAD [13] repository. This short paper presents the initial findings of our study.

Our key finding is that most Wi-Fi networks are significantly underutilized. The median medium utilization is less than 40% in all settings that we study even during the busiest times and much lower during other times. To put this number in perspective, consider that even a single saturated 802.11 transmitter can occupy over 70% of the medium [7].

The common-case of low utilization renders many active research problems less urgent, at least until Wi-Fi networks become heavily utilized again. We discuss these in detail in the body of the paper. But briefly, based on our findings, we argue that problems such as rate anomaly [20], chaos due to the presence of multiple, overlapping but independent networks [2], hidden and exposed terminals [18], and efficient network coding [26] are less pressing. We do not claim that these problems do not merit any attention, but it is likely that simpler and perhaps less effective solutions would suffice at present.

At the same time, we argue that certain other problems need rethinking in the light of low utilization. For instance, more effective autorate algorithms and loss protection schemes can be re-designed to take advantage of the spare capacity. Other aspects of wireless networks that merit renewed attention are the analytical models of MAC behavior and experimental workloads, both of which are commonly driven today by a world view of heavy utilization [33].

2. METHODOLOGY

The results presented in the paper are based on packet traces that we collected from a variety of locations, as well as traces from the CRAWDAD repository. In this section, we describe our monitoring equipment, and the data sets.

2.1 Measurement Setup

Our measurement setup is common to all locations. Depending on the size of the location, we deploy a number of sniffer, dispersed throughout the location. The nodes are small form-factor PCs equipped with Atheros cards and run Windows XP. The sniffers are configured to be in RFMON mode and record frames at the MAC layer, including the prism header of each frame. We modified the drivers to also record packets that fail the CRC check. The sniffers were recording packets on 802.11g channels 1, 6 and 11 and were switching channels every 20 minutes. The multiple sniffers at the same location were loosely synchronized to change channels at the same time. The sniffers generate a small amount of probe traffic, but are otherwise passive.

2.2 Data Sets

We collected data several Wi-Fi networks. We categorize these networks into five categories, as shown in Table 1.

Apartments in multi-unit buildings: We deployed sniffers in six residential apartments and UCSB student dormitories. These locations tend to overhear a lot of traffic from competing networks.

Single family houses: Single family houses do not share walls with other units, and typically overhear little wireless activity from neighboring networks. We collected data from three single-family homes in the Seattle area.

Enterprise Networks: Enterprise networks typically have well planned wireless deployments. We collected measure-

Location	Instances(#)	Sniffers(#)	Busy Time
Apartments/Dorms	7	6	18:00 - 00:00
Single family houses	3	10	18:00 - 22:00
Enterprise Networks.	2 1‡	5 192	9:00 - 18:00 9:00 - 18:00
Large gathering	1‡	8	9:00 - 18:00
Coffee Shop	2	3	10:00 - 19:00

Table 1: Data sets used in this study. Data sets marked with ‡ are from the CRAWDAD repository.

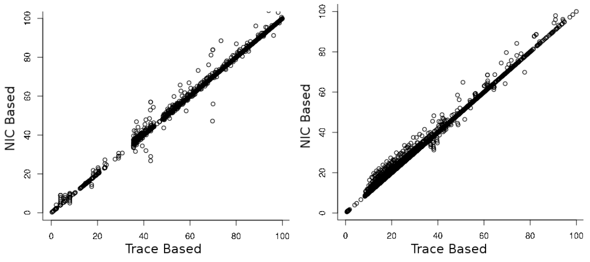

(a) Dormitory (b) Coffee Shop

Figure 1: Comparison of NIC-based and trace-based utilization

ments from within one of the Microsoft office buildings as well as the CS department at U. of Cal at Santa Barbara. In addition, we analyzed CRAWDAD trace representing the wireless activity in the UCSD CS building collected on Thursday, January 11, 2007.

Large Gathering: We analyzed wireless traces collected at the IETF meeting in San Diego, CA in November 2006. Such gatherings typically have a large number of wireless users. This trace is from the CRAWDAD repository.

Coffee shop/hotspot: We deployed sniffers in a popular coffee shop in downtown Santa Barbara. For each of the two datasets, we collected measurements for one full day on a weekend, and the coffee shop had about 50 users at the peak.

At each location, we collected measurements for at least 24 hours in each except at the coffee shop where we collected during the cafe hours. This amounts to a total of over 350 hours. However, for the purpose of this study, we only focus on peak usage times. For example, for enterprise networks, we only use data from during normal office hours. This ensures that our analysis is not skewed by periods during which the networks are unlikely to in use (e.g. late nights at corporate offices). The specific time periods were chosen based on the results from prior analysis of DSL and IP networks [16, 5], as well as our own inspection of the traces.

2.3 Computing Medium Utilization

We use the term 'utilization' of a wireless network to denote the fraction of time that the medium was utilized in some time interval as a result of packet transmissions or other interfering traffic in the same part of the spectrum. Hence each node has its own view of medium utilization depending on the activity in its vicinity.

The airtime utilization at a node can be computed in two ways, either using packet traces [23, 34], or by using low-level information from the NIC [1].

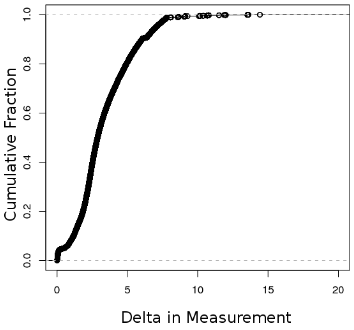

Figure 2: Difference between NIC-based and trace-based utilization

Trace-based utilization: To compute airtime utilization using packet traces, we use the same methodology as [23, 34]. We consider all packets, including those that failed CRC checks. We compute the transmission time of a packet using its size and transmission rate. For each packet we add the appropriate overheads such as preamble and IFS intervals. Since we do not have information on the type of preamble used at the MAC layer, we use the long preamble in all our computations, which errs on the side of overestimating the medium utilization. This overestimation explains the near 100% utilization that we see at times.

However, this method can underestimate utilization in some cases. Packets dropped due to preamble corruption are never recorded by the sniffer and hence not included in this computation. Noise due to non 802.11 devices such as microwaves and phones is also unaccounted.

NIC-based utilization: The Atheros-based NICs continuously monitor the energy in the spectrum to determine whether the channel is free and maintain statistics on how often the medium was sensed free or busy. We access these statistics using driver hooks [1]. This information provides a more complete view of medium utilization because the NIC senses the medium as busy even if a packet is discarded by the PHY or MAC layer and even when the energy in the medium is from non-802.11 devices.

However, if a large majority of packets are received with very low signal strength (less than than the CCA threshold), the reported utilization may be an underestimate [1].

In Figure 1(a) and (b), we compare the NIC-based and trace-based utilization for two locations: dormitory and coffee shop. Each point is a different 1 second interval. We find that the medium utilizations computed by the two methods are quite similar. Figure 2 shows the CDF of the difference in utilization values computed using the two methods. We see that the utilization values differ by less than 10% over 90% of the times.

Since the difference in the two estimations is small, we present results from packet trace data in the rest of the paper. This lets us study CRAWDAD traces in a comparable manner, since we do not have NIC-level medium utilization information for these traces. airtime as reported by the wire-

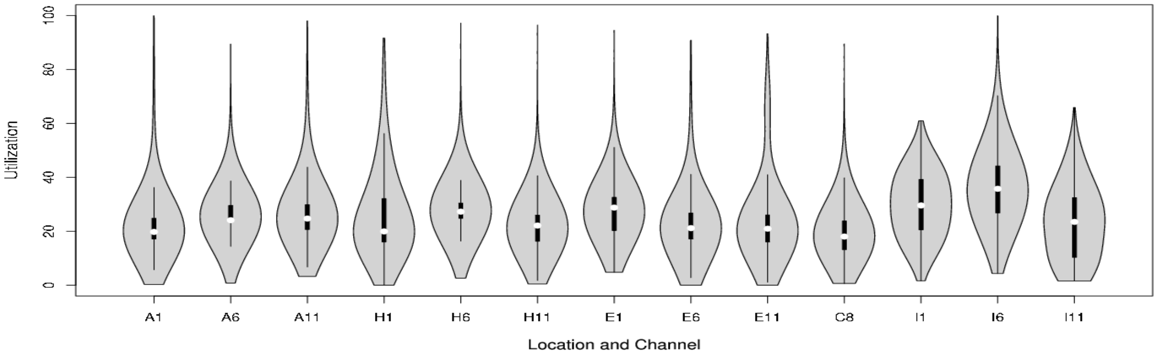

Figure 3: Violin chart that showing the quartile distribution of utilization values.

less card as well as captured packets.

3. WIRELESS NETWORK UTILIZATION

We study the overall medium utilization in each of the datasets. In Figure 3 is a violin plot [37] of the medium utilization observed from the packet traces.

The violin plot is a combination of the box-plot and a probability density function. The box-plot represents the median, upper and lower quartiles. The whiskers extend to 1.5 times the inter-quartile range. Around the box-plot is the probability density curve of different utilization values. In short, the width of the violin indicates the probability density of the corresponding y-axis value.

Each violin represents the aggregated results from all sniffers on a particular channel for each dataset. Thus, A1 represents aggregated results from all sniffers that were placed in apartments and dorms on channel 1 and A6 and A11 are from channels 6 and 11 respectively. Similarly, we have 3 violins for each of the other datasets: houses (H1-H11), enterprise traces (E1-E11), UCSD traces (U1-U11) and IETF traces (I1-I11) and one for the coffee shop where the AP was on channel 8 (C8).

The utilization was measured at 1 second granularity, and for each location where there were multiple sniffers, we take the maximum utilization reported to get an upper bound on that location. For example, we had deployed two sniffers in the coffee shop. For each 1 second interval, we picked the sniffer that reported the larger airtime utilization.

We see that except for some IETF traces, the median utilization is less than 30%, which is quite low. Even the 75th percentile is less than 40% in most cases. The shape of the violin indicates that the maximum density lies in the median range. The shape also indicates that while we do see very high values of utilization, this happens quite rarely.

These observations beg the question: why is the utilization low? We try to answer this question next.

3.1 Why underutilized?

There are several potential reason for the observed underutilization of these networks. It could be that there were not enough APs or clients in the area we monitored. Or perhaps the loss rate was so high that higher layer protocols like TCP backed off, or the users simply found the performance unacceptable. Or, it could simply be that there was just not enough demand. We analyze these possibilities one by one.

3.1.1 Number of APs

Several recent research studies [2] have posited that the dense and unplanned nature of wireless network deployments can lead to significant performance degradation. The distribution of number of distinct APs seen in our traces during distinct 1-second interval is shown in Figure 4(a). The median number of APs we see is about 4. This implies that we see a moderate number of networks.

In Figure 4(b), we plot the utilization as a function of number of APs seen. We group the 1-second intervals according to the number of APs observed in that interval. For each group, the graph shows the median, and 10^{th} and 90^{th} percentile of the utilization.

We see that there is no correlation between the number of APs seen, and the utilization. The median is always around 30%, regardless of the number of APs. Thus, the low utilization we observe is probably not because there aren't enough Wi-Fi networks in the areas we monitored. Could it be because there were not enough clients? We look at this next.

3.1.2 Number of active clients

The distribution of number of distinct clients seen in our traces during 1-second interval is shown in Figure 5(a). The median number of clients we see is around 10. Figure 5(b) shows that there is little correlation between the number of clients and the utilization: even when we see 20 clients, the median utilization is still less than 40% [1].

Thus, the low utilization we observe is not because there are not enough clients in the areas we monitored. Could it be that these clients suffered from heavy losses, and thus higher-layer protocol backed off?

3.1.3 Loss Rate

Determining correlation between utilization and loss rate is quite difficult. It requires detailed monitoring and complex analysis [29, 11] to *passively* determine the loss rate

---

[1] Some of these 1 second intervals are from the IETF trace, where we see a lot of clients, announcing presence null packets.

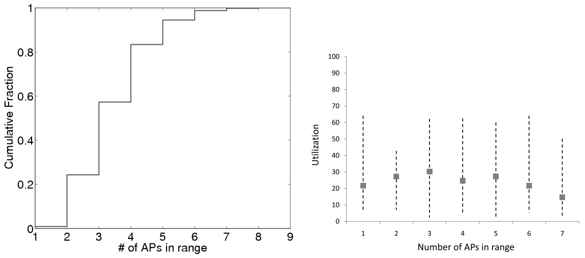

(a) CDF of Number of APs (b) APs vs. Utilization

Figure 4: Analysis of number of APs in range

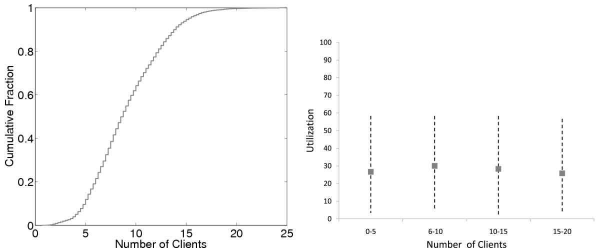

(a) CDF of Number of Clients (b) Clients vs. Utilization

Figure 5: Analysis of number of clients in range

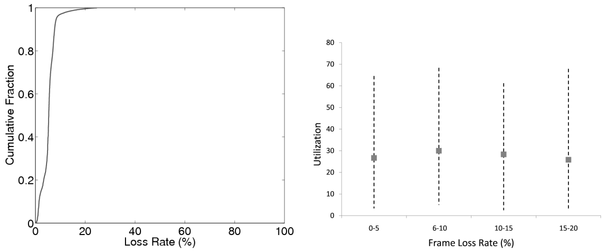

(a) CDF of loss rates (b) Loss vs. Utilization

Figure 6: Analysis of loss rate

the clients are experiencing. Instead, we approximate the loss rate by the fraction of packets that had the retransmit bit set. While this method underestimates the loss rate, [29] has shown that the error is less than 20% in 85% of the time.

The CDF of loss rate is shown in Figure 6(a). We see that the loss rate is less than 10% 90% of the time. These numbers are similar to those observed in [11]. Furthermore, Figure 6(b) shows that there is no correlation between loss rate and medium utilization.

Notice that the loss rates we compute above are PHY-layer loss rates. The ARQ mechanism used in 802.11 MAC can be quite effective at lowering the loss rates seen by higher layers. For example, a 20% loss rate at PHY layer translates to only 0.16% loss rate above the link layer after just four retransmissions.

Thus, it appears that high loss rate is not a cause for low utilization that we have observed. Still, we plan to study the loss characteristics in more detail (e.g. the burstiness) to arrive at a definite conclusion. For example, in [11], authors have shown that wireless losses are responsible for limiting the throughput of about 20% of the TCP transactions.

The remaining (and likely the simplest) explanation is that average utilization is low because the average demand is low. This is true even of corporate wireline networks, as data presented in [15] indicates. We plan to study our trace data in more detail to shed more light on this possibility. If the low demand is indeed the cause for low utilization, then it should be possible to put the remaining capacity to good use (e.g. to reduce loss rate even further by aggressive use of FEC).

Of course, any scheme to use available capacity to hide losses or improve other aspects of wireless performance is feasible only if periods of low utilization are available for some significant length of time. If periods of low utilization have very short "run lengths", it may not be possible to take full advantage of them. To this end, we study the distribution of duration of periods of low utilization.

3.2 Distribution of periods of low utilization

To gain insight into how the periods of low utilization are distributed, we model the system using a simple two-state transition model. The time is divided into 1 second intervals. The system is said to be in a low state when the utilization is under 50% and in a high state when the utilization is over 50%. We compute the four transition probabilities for this system using the data sets plotted in Figure 3. The resulting model is shown in Figure 7(a). The model says that the probability of the system being in low utilization state is 0.75. Furthermore, once the utilization is low, it stays low with probability of 0.77, which is quite high.

Another way to look at the same data is to plot a CDF of "run lengths" of low and high utilization periods, as shown in Figure 7(b). We see that long periods of low utilization periods are common, and most of them are around 40 minutes in length. On the other hand, the median length of a busy period is about 10 minutes. These results indicate it might be feasible to exploit periods of low utilization to improve Wi-Fi performance.

3.3 Generality

Like any measurement-based study, our results are based on a finite set of measurements. Nonetheless, we believe that our key observation, that wireless networks are underutilized is generally applicable. While we can not provide a quantitative proof of generality, we offer several arguments in support of our claim. First, we have analyzed data for several diverse scenarios, such as homes and offices, where wireless networks are commonly used . Second, the data has been filtered to include only the most active time segments. The airtime utilization in other time periods is even lower. Third, when computing utilization for a given time period in a given network, we always report the largest utilization from among all the monitoring nodes. Fourth, we analyze publicly available traces gathered by other researchers, and they continue to support our conclusions as well.

4. IMPLICATIONS

The measurements in the previous section have shown that

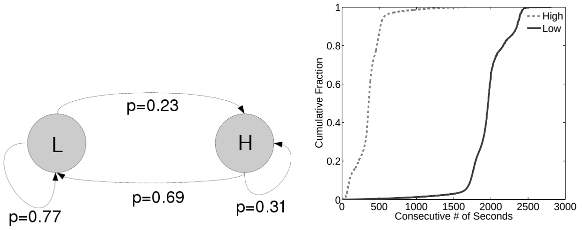

(a) Transition probabilities (b) Distribution of run lengths

Figure 7: Distribution of low utilization periods

typically there is plenty of free airtime in the vicinity of WiFi access points, even during peak usage times. If this low utilization scenarios were to persist in the future, it has implications as to which active research problems are or are not pressing and which problems require rethinking.

4.1 Problems that are less pressing

In this section, we list five problems that we deem less pressing, at least in the near future. *We do not claim that researchers should stop looking at these problems altogether.* But because the presence of ample free airtime makes these problems less severe, there might be simpler (but perhaps not as effective) mechanisms that bring most of the gain.

(i) Rate anomaly: A problem related to the availability of multiple transmission rates in 802.11 that has been the focus of much work is the rate anomaly problem [35]. This problem is inherent in the design of the 802.11 MAC. The MAC arbitrates channel access at the granularity of packets. So, a sender that uses lower transmission rate consumes more air time, and can impact the performance of other senders with higher transmission rate. Many solutions to this problem have been proposed [3, 14]. Most require changes to the MAC protocol itself.

However, it is important to note that this problem occurs *only* when the medium is heavily utilized. If there is plenty of free airtime, disparate transmission rates cannot cause a problem. Our data shows that free airtime is plentiful in today's wireless networks. Thus, we believe that rate anomaly is not a serious performance issue.

(ii) Network Chaos: Researchers have recently began to ponder the chaotic nature of WiFi deployments [2]. In many environments (e.g. urban downtowns, multi-family residential buildings) several WiFi networks are deployed independently. There is no central controller to perform coordinated channel and power assignment. Researchers have worried that this can lead to inefficiencies and poor performance for everyone. Several complex game theoretic strategies have been proposed to tackle the problem.

However, in most cases, the chaotic nature of these deployments is a problem only if the airtime utilization is high. The mere presence of a large number of APs/wireless networks does not cause any problems, regardless of what channels they operate on. This factor stands out in our data, as we find no correlation of free airtime with and number of APs present and plenty of free airtime is available even in the presence of six or seven APs. Thus, the need for complex, game-theoretic channel assignment algorithms is unclear.

(iii) Hidden Terminals: The hidden terminal problem is well known [18, 31]. It occurs when transmissions from two nodes, that can not hear each other, collide at the receiver for one of the nodes. Numerous solutions to solve this problem have been proposed [6, 25] to address the problem. The 802.11 standard recommends the use of RTS/CTS exchange to avoid hidden terminal problems.

However, such collisions will be cause significant performance problems only if the network utilization is high, or the two senders somehow get synchronized with each other. The latter is unlikely to happen in practice. In underutilized networks, the hidden terminals will cause only occasional collisions (since the senders are not sending much data) and the resulting loss can be easily handled by MAC-level backoff and retransmissions mechanisms with little impact on user-perceived performance. We believe that this one of the reasons that hidden terminals is not reported as a major performance problem in deployed networks, even though most networks do not use RTS/CTS due to overhead it imposes.

(iv) Exposed Terminals: The exposed terminal problem occurs when a node is prevented form sending due to the presence of another transmitter nearby. This occurs because the carrier sense mechanism used in 802.11 is conservative, and prevents a node from transmitting when another node is transmitting, for the fear of causing a collision. Several solution to this problem have been proposed as well [38]. Most require modifying the carrier sense mechanism in some way.

However, significant throughput reduction due to exposed terminals is a problem only of network utilization is very high. Otherwise, both nodes will have ample opportunity to send their packets. Thus, we believe that exposed terminals is not a significant problem in today's networks.

(v) More capacity: (802.11n, network coding): If today's wireless networks are mostly undersubscribed, the rationale for upgrading to higher-throughput standards such as 802.11n becomes weaker. However, the MIMO technology of the 802.11n PHY layer not only improves throughput but can also reduce loss and increase coverage. These other two aspects are still important. Better coverage or the ability to use fewer APs is easier to justify.

However, the need for certain other advances in PHY layer technologies is more questionable. A large body of recent research in wireless networking is focused on network coding, and complex PHY layer symbol manipulations to enable wireless nodes exchange more packets with fewer transmissions [18]. Our data indicate that these techniques are not urgently needed for today's wireless networks.

While our studies look at the average usage of average scenarios, we acknowledge that networks are likely to experience periods of high utilization, during which the problems that are "less pressing" will be very important. However, the results show that these situations are not the common case. Also, they build a case for capacity-aware design for several problems, which we outline in Section 4.2

4.2 Problems that need rethinking

In this section, we list four problems that merit rethinking in a world where the plenty of air time is available.

(i) Autorate Algorithms: Autorate algorithms control which of the many available transmission rates a sender uses. The goal is to balance the tradeoff between loss rate and air time consumption. Often (but not always), lower transmission rate results in lower loss rates, but it (always) leads to increased airtime utilization. However, current autorate algorithms do not consider the available free airtime when making their decisions and instead aim for efficiency (e.g., minimize airtime usage per successfully transmitted bit) [9, 21].

We propose that autorate algorithms should consider sender's backlog and the airtime utilization while making rate adjustment decisions. For example, if the sender is not backlogged, and the airtime utilization is low, why increase the transmission rate? It will not increase sender's throughput (since there is no backlog), nor will it improve anyone else's throughput (low air time utilization). On the other hand, it may increase sender's loss rate. Needless to say, many other details (e.g. impact on jitter, impact on faraway stations) must be considered while designing a detailed scheme.

(ii) Loss protection: The two most important factors that determine the performance of a wireless client is the amount of free medium and the loss rate of its transmissions. We find that while most clients today are not limited by medium availability, losses continue to be a problem. The current 802.11 standard uses ARQ (i.e., retransmissions) for loss recovery. While ARQ-based loss recovery is efficient in terms of bandwidth consumed, it introduces extra delay and jitter.

If the efficiency in terms of airtime usage is not a prime concern, (which our data show that it should not be in many situations), more aggressive loss protections mechanisms become desirable. One possibility is to proactively add more FEC (forward error correction) bits. Another way is to employ lower transmission rates that consume more airtime but add more redundancy for higher loss protection. The presence of free air time and the occurrence of loss even at low 802.11 transmission rates makes the case for developing even lower transmission rates. Such rates will also improve wireless coverage, which can be a problem in some settings [10].

(iii) Analytical Models: Analytical models of MAC behavior are valuable in understanding wireless performance as well as planning. However, most models consider the case of saturated medium with backlogged senders [40, 8, 39]. However, we see that the common case is that of an unsaturated medium. Modeling this scenario is more involved because the designers need to factor in how much and when a node might transmit. Recent work in the domain of wireless meshes has considered the unsaturated setting [28]. Similar models are needed for infrastructure networks.

(iv) Realistic Experimental Workloads: It it not only the current analytical work that is not capturing the common case of low utilization; current experimental works has this shortcoming as well. Perhaps driven by a vision of over-subscribed wireless network, most experimental studies of new wireless technologies focus on scenarios in which transmitters are saturated and send data non-stop [30]. To better understand the value of proposed enhancements, we should develop models of workload observed in real environments and use these models to guide experiments.

5. RELATED WORK

As we mentioned earlier, many researchers have studied various properties of wireless networks, in a wide variety of settings. Here, we briefly describe a few of these studies.

Home Networks: [32] studies the wireless network performance in home environments. This study provides early evidence of significant variability and asymmetry in home network link quality. Using a testbed deployed at three homes, this work measured the TCP and UDP throughput obtained, evaluation of the impact of automatic rate selection, and comparison of the impact of flexible topologies on the performance of home wireless networks.

Conference Networks: [24, 23] study the IETF network and characterizes the high utilization and loss rates that exist in these networks. [34] collected link layer traces from the SIGCOMM conference and analyzed the causes for high retransmission rate in the network .

Campus Networks: Several studies have focused on analysis of wireless network usage in campuses [36, 27], including application workloads and session durations [4, 19]. These studies were based on the analysis of wired distribution network traffic and polled SNMP management data.

Wit [29] and Jigsaw [12] are built to understand how 802.11 networks have in their full empirical complexity. They present robust merging procedures to combine the necessarily incomplete views from multiple, independent monitors into a single, more complete trace of wireless activity to obtain a cross-layer viewpoint to isolate performance artifacts.

Public Hotspots: Giroire et. al. have studied how the wireless usage patterns in office and coffee shop environments differ from each other.

In this paper, we study a critical yet unexplored aspect of wireless networks: the extent to which commonly deployed wireless networks are utilized. We differ from the earlier work in that we do not look at *how* users use the network, but we study *how much* of the network capacity is used.

6. CONCLUSIONS

In this paper, we address the question as to how congested today's wireless networks are. We present measurements from several diverse wireless networks using both data that we collected as well as publicly available data sets. The analysis confirms that wireless networks have plenty of free capacity, even during times of peak usage. We conclude with a discussion on several specific research problems such as autorate algorithms and loss recovery that could be redesigned given that plenty of airtime is available. Also several problems such as complex channel assignment schemes become less urgent.

7. REFERENCES

[1] P. Acharya, A. Sharma, E. Belding, K. Almeroth, and K. Papagiannaki. Congestion-Aware Rate Adaptation in Wireless Networks: A Measurement-Driven Approach. In *Proc of SECON*, 2008.

[2] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self-management in Chaotic Wireless Deployments. In *Proc of MOBICOM*, 2005.

[3] V. Bahl, R. Chandra, P. P. C. Lee, V. Mishra, J. Padhye, D. Rubenstein, and Y. Yu. Opportunistic Use of Client Repeaters to Improve Performance of WLANs . In *Proc. of CONEXT*, 2008.

[4] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan. Characterizing user behavior and network performance in a public wireless LAN. In *Proc of ACM SIGMETRICS*, 2002.

[5] M. Balazinska and P. Castro. Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network. In *Proc. of MobiSys*, 2003.

[6] V. Bharghavan, A. J. Demers, S. Shenker, and L. Zhang. MACAW: A Media Access Protocol for Wireless LANs. In *SIGCOMM*, 1994.

[7] G. Bianchi. IEEE 802.11-saturation Throughput Analysis. *Communications Letters, IEEE*, 2(12):318–320, 1998.

[8] G. Bianchi. Performance Analysis of the IEEE 802.11 Distributed Coordination Function. *IEEE JSAC*, 18(3), 2000.

[9] J. Bicket. Bit-rate Selection in Wireless Networks. *Master's Thesis, Massachusets Institute of Technology*, 2006.

[10] R. Chandra, J. Padhye, A. Wolman, and B. Zill. A Location-based Management System for Enterprise Wireless LANs. In *Proc. of NSDI*, 2007.

[11] Y.-C. Cheng, M. Afanasyev, P. Verkaik, P. Benko, J. Chiang, A. C. Snoeren, S. Savage, and G. M. Voelker. Automating cross-layer diagnosis of enterprise wireless networks. In *Proc. of SIGCOMM*, 2007.

[12] Y.-C. Cheng, J. Bellardo, P. Benko, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: Solving the puzzle of enterprise 802.11 analysis. In *Proc. of SIGCOMM*, 2006.

[13] A Community Resource for Archiving Wireless Data at Dartmouth. http://crawdad.cs.dartmouth.edu/.

[14] J. DUNN, M. NEUFELD, A. SHETH, D. GRUNWALD, and J. BENNETT. A practical cross-layer mechanism for fairness in 802.11 networks. *MONET*, 11(1):37–45, 2006.

[15] J. Eriksson, S. Agarwal, P. Bahl, and J. Padhye. Feasibility Study of Mesh Networks for All-Wireless Offices. In *Proc. of MOBISYS*, 2006.

[16] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True. Deriving Traffic Demands for Operational IP Networks: Methodology and Experience. *IEEE/ACM Transactions on Networking*, 2001.

[17] F. Giroire1, J. Chandrashekar, G. Iannaccone, K. Papagiannaki, E. M. Schooler, , and N. Taft. The cubicle vs. the coffee shop: Behavioral modes in enterprise end-users. In *Proc. of PAM*, 2008.

[18] S. Gollakota and D. Katabi. Zigzag Decoding: Combating Hidden Terminals in Wireless Networks . In *Proc. of SIGCOMM*, 2008.

[19] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of the ACM MobiCom*, pages 187–201, Philadelphia, PA, Sept. 2004.

[20] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance Anomaly of 802.11b. In *Proc. of INFOCOM*, 2003.

[21] G. Holland, N. Vaidya, and P. Bahl. A Rate-adaptive MAC Protocol for Multi-Hop Wireless Networks. In *Proc. of MOBICOM*, 2001.

[22] R. Jain, D. Lelescu, and M. Balakrishnan. Model T: An Empirical Model for User Registration Patterns in a Campus Wireless LAN. In *mobicom*, pages 170–184, 2005.

[23] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding congestion in ieee 802.11b wireless networks. In *Proceedings of IMC*, Oct. 2005.

[24] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding Link-Layer Behavior in Highly Congested IEEE 802.11b Wireless Networks. In *Proceedings of EWIND*, pages 11–16, Philadelphia, PA, Aug. 2005.

[25] P. Karn. MACA – A New Channel Access Method for Packet Radio. In *Proc. of ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, 1990.

[26] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, and J. Crowcroft. XORs in the Air: Practical Wireless Network Coding. In *Proc. of SIGCOMM*, 2006.

[27] D. Kotz and K. Essien. Analysis of a campus-wide wireless network. In *Proceedings of the ACM MobiCom*, Sept. 2002.

[28] Y. Li, L. Qiu, Y. Zhang, R. Mahajan, and E. Rozner. Predictable Performance Optimization for Wireless Networks. In *Proc. of SIGCOMM*, 2008.

[29] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Analyzing the mac-level behavior of wireless networks in the wild. In *Proc. of SIGCOMM*, 2006.

[30] R. Murty, J. Padhye, R. Chandra, A. Wolman, and B. Zill. Designing High Performance Enterprise Wi-Fi Networks . In *Proc. of NSDI*, 2008.

[31] P. C. Ng, S. C. Liew, K. C. Sha, , and W. T. To. Experimental study of hidden node problem in ieee 802.11 wireless networks. In *SICOMM poster*, 2005.

[32] K. Papagiannaki, M. Yarvis, and W. S. Conner. Experimental Characterization of Home Wireless Networks and Design Implications. In *Proc. of INFOCOM*, 2006.

[33] L. Qiu, Y. Zhang, F. Wang, M. K. Han, and R. Mahajan. A General Model of Wireless Interference. In *Proc. of MOBICOM*, 2007.

[34] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan. Measurement-based characterization of 802.11 in a hotspot setting. In *Proceedings of EWIND*, pages 5–10, Aug. 2005.

[35] G. Tan and J. Guttag. Time-based fairness improves performance in multi-rate wlans. In *Proceedings of USENIX*, 2004.

[36] D. Tang and M. Baker. Analysis of a local-area wireless network. In *Proceedings of the ACM MobiCom*, pages 1–10, Boston, MA, Aug. 2000.

[37] Package vioplot.

[38] M. Vutukuru, K. Jamieson, and H. Balakrishnan. Harnessing Exposed Terminals in Wireless Networks. In *Proc of NSDI*, 2008.

[39] H. Wu, Y. Peng, K. Long, S. Cheng, and J. Ma. Performance of reliable transport protocol over IEEE 802.11 wireless lan: analysis and enhancement. In *INFOCOM*, 2002.

[40] H. Zhai, Y. Kwon, and Y. Fang. Performance analysis of ieee 802.11 mac protocols in wireless lans: Research articles. *Wireless Communication and Mobile Computing*, 4(8), 2004.