

Ambient Interference Effects in Wi-Fi Networks

Aniket Mahanti¹, Niklas Carlsson¹, Carey Williamson¹, and Martin Arlitt^{1,2}

¹ Department of Computer Science, University of Calgary, Calgary, Canada
{amahanti,ncarlsson,carey}@cpsc.ucalgary.ca

² Sustainable IT Ecosystem Lab, HP Labs, Palo Alto, U.S.A.
martin.arlitt@hp.com

Abstract. This paper presents a measurement study of interference from six common devices that use the same 2.4 GHz ISM band as the IEEE 802.11 protocol. Using both controlled experiments and production environment measurements, we quantify the impact of these devices on the performance of 802.11 Wi-Fi networks. In our controlled experiments, we characterize the interference properties of these devices, as well as measure and discuss implications of interference on data, video, and voice traffic. Finally, we use measurements from a campus network to understand the impact of interference on the operational performance of the network. Overall, we find that the campus network is exposed to a large variety of non-Wi-Fi devices, and that these devices can have a significant impact on the interference level in the network.

Keywords: Wi-Fi, Interference, Spectrogram, Duty Cycle, Data, Video, Voice.

1 Introduction

Wireless Fidelity (Wi-Fi) networks allow users with wireless-capable devices to access the Internet without being physically tied to a specific location. Today's Wi-Fi networks are not only being used for Web surfing, but also for viewing live video content (e.g., live sports and news events) and voice communication (e.g., voice over IP or VoIP). Many of these services require high quality of service and reliability, which can be degraded by wireless interference.

Wi-Fi networks employ the IEEE 802.11 protocol, which uses the unlicensed 2.4 GHz Industrial, Scientific, and Medical (ISM) Radio Frequency (RF) band [6]. Since the ISM band is unlicensed, it is available for use by multiple devices (both Wi-Fi and non-Wi-Fi), inherently causing interference for one another. The 802.11 protocol is considered to be a polite protocol in that an 802.11 device will transmit only if it senses that the RF channel is free. Non-Wi-Fi devices such as microwave ovens are oblivious to this protocol. These devices transmit regardless of whether the channel is free or not. This makes the interference problem challenging in Wi-Fi networks.

Network practitioners often deploy Wi-Fi networks without knowledge of the ambient usage of the ISM band by non-Wi-Fi devices. For example, trace capture programs,

* ©IFIP, (2010). This is the author's version of the work. It is posted here by permission of IFIP for your personal use. Not for redistribution. IFIP Networking 2010, Chennai, India.

employed for site surveys used to select RF channels for problem-free functioning of Wi-Fi networks, only recognize devices using the 802.11 protocol. Since this process concentrates on the link layer and up, any activity from non-Wi-Fi devices is ignored. Unfortunately, since several non-Wi-Fi devices may be operating on the same channel as a Wi-Fi network, and cause severe interference, it is difficult to identify interference sources using this trace capture technique.

For a more complete picture of the activity on the 2.4 GHz ISM band, one needs to consider the physical layer. In this paper, we use an off-the-shelf wireless spectrum analyzer to understand how non-Wi-Fi devices impact the functioning of Wi-Fi networks. Using controlled experiments, we characterize the interference properties of six non-802.11 devices. Five are *unintentional* interferers: a microwave oven, two cordless phones (one analog and one digital), an analog wireless camera, and a Bluetooth headset. We also evaluate one *intentional* interferer, a wireless jammer, for comparison purposes. In addition to capturing the basic characteristics of these devices, we measure, quantify, and discuss implications of their interference on data, video, and voice traffic. Finally, using passive measurements from an operational campus network, we try to understand the impact of interference on the network.

Our results show that among the unintentional interferers, microwave ovens, analog cordless phones, and wireless cameras have the most adverse impact on Wi-Fi networks. Because of its wideband interference, microwave signals affect several Wi-Fi channels at close range, however, their impact is still felt at longer distances. Analog cordless phones and wireless cameras are continuous narrowband interferers that completely obliterate Wi-Fi service on any channels they are using. Digital cordless phones and Bluetooth headsets have minimal impact on Wi-Fi because of their frequency hopping nature. From our production network measurements, we find that the campus network is exposed to a large variety of non-Wi-Fi devices, and that these devices can have a significant impact on the interference level in the network. For example, during certain times of the day, almost 80% of a channel may be occupied by interferers, and it is common to see some interference device active (in the background) almost all the time.

The rest of the paper is organized as follows. Section 2 discusses the interferers we studied and our experimental setup. Section 3 describes the organization of the Wi-Fi channels and characterizes the physical layer properties of the six interferers. Section 4 studies the impact of non-802.11 interference on data, video, and voice traffic over Wi-Fi. Section 5 studies the channel utilization of Wi-Fi and various interferers in a production network. Section 6 describes related work. Section 7 concludes the paper.

2 Methodology

2.1 Interferers

Microwave Oven: We used a Panasonic NNS615W microwave oven. This device had a maximum output power of 1,200 W. During our experiments, the oven was operating at maximum power and had a container of food inside the oven.

Analog Wireless Video Camera: We used a Lorex SG8840F analog wireless video camera operating at 2.4 GHz. These cameras provide long-range surveillance using analog signals, typically use directional antennas, and can have a range up to 1.5 km.

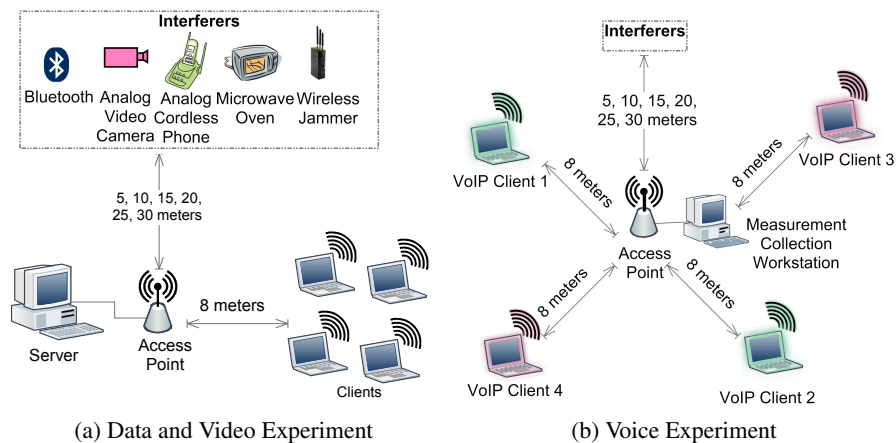


Fig. 1. Experimental setup

Analog Cordless Phone: We used a Vtech GZ2456 analog cordless phone operating at 2.4 GHz. The phone consists of a base and a handset, and as with other analog phones, works by converting voice into electronic pulses and sending them on a designated channel between the base and handset.

Digital Cordless Phone: We used a Uniden DCT648 digital phone (base and handset) operating at 2.4 GHz. This particular phone uses Digital Spread Spectrum (DSS) to change channels frequently for enhanced voice quality and security.

Bluetooth Headset: We used a Plantronics Pulsar 590A headset. These are short range devices and have a maximum operating range of 10 m.

Wireless Jammer: Wireless jammers may be used for RF Denial of Service attacks. The power levels of these jammers vary from 1 mW to 30 mW.

2.2 Experimental Setup

We conducted two sets of measurement experiments to study the physical layer characteristics of the above-mentioned interferers and their impact on Wi-Fi networks.

The first set of experiments examines the physical-layer characteristics of these interferers in an isolated environment. We used an off-the-shelf spectrum analyzer called AirMagnet Spectrum Analyzer³ for this purpose. The spectrum analyzer is a hardware/software unit that consists of a radio for detecting RF energy in the 2.4 GHz ISM band, and a software engine that performs Fast Fourier Transforms (FFTs). The spectrum analyzer uses these FFTs to classify known interferers. The physical layer measurements were taken in an interference-neutral environment such that we only captured RF energy from the specific device under study.

The second set of experiments quantifies the impact of interferers on the performance of a Wi-Fi network. We measured the performance degradation of the network

³ http://www.airmagnet.com/products/spectrum_analyzer/




Fig. 2. Structure of Wi-Fi Channels in the 2.4 GHz band

in the presence of the interferers for three types of traffic workloads: data, video, and voice. Our Wi-Fi network consisted of a single D-Link 2100 access point (AP) running in IEEE 802.11g mode. We chose the 802.11g standard because it offers higher transmission rates than 802.11b and is a popular choice for deployment in enterprise, academic, home, and hot spot networks. Additionally, most modern laptops are equipped with 802.11g capable Network Interface Cards (NICs).

We used the setup shown in Figure 1(a) for data and video experiments. We connected a server workstation using an Ethernet cable. We placed four Lenovo T61 laptop clients with built-in 802.11abg NICs at a distance of 8 m from the AP. We placed a single interferer at successive distances of 5, 10, 15, 20, 25, and 30 m from the AP to record the performance degradation of the interferer on the network. For each interferer, we repeated this scenario four times to record our readings.

The experimental setup for the voice experiment is shown in Figure 1(b). The AP was placed in the centre and was connected to a measurement workstation using an Ethernet cable. We placed two pairs of laptops facing each other diagonally at a distance of 8 m from the AP. We configured Ekiga⁴ VoIP software on the laptops to allow direct IP calling without using a SIP server. VoIP client 1 was communicating with VoIP client 2 and VoIP client 3 was communicating with VoIP client 4. Interferer placement was the same as in the case of the data and video experiments.

3 Physical-Layer Characteristics

3.1 Channel Structure in Wi-Fi

Wi-Fi channels are organized into 14 overlapping channels each having a spectral bandwidth of 22 MHz. Figure 2 shows a graphical representation of the Wi-Fi channels in the 2.4 GHz band. The figure shows the centre frequencies of each Wi-Fi channel. Adjacent channels are separated by 5 MHz, except for channel 14 whose centre frequency is separated from channel 13 by 12 MHz. A single channel can handle 50 or more simultaneous users [6]. Usage of Wi-Fi channels are governed by national regulatory

⁴ <http://ekiga.org/>

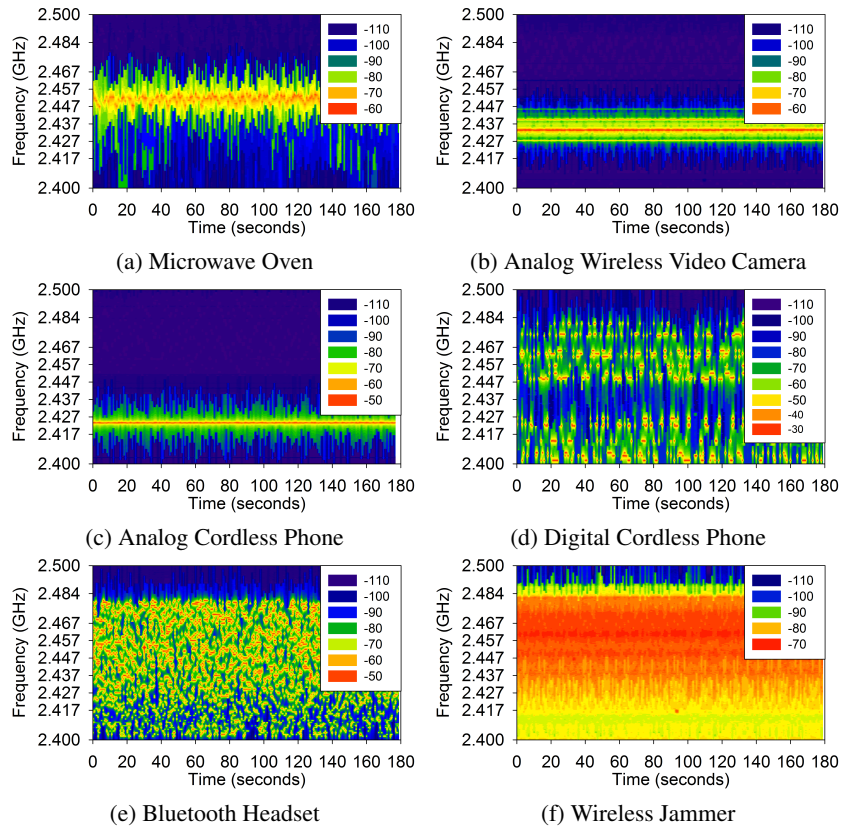


Fig. 3. Spectrograms for six non-Wi-Fi devices

agencies of the respective countries. In North America, only the first 11 channels are available for use. In the rest of the world the first 13 channels are available for use. Japan allows the use of channel 14 as well, however, it is only available to 802.11b using Direct Sequence Spread Spectrum (DSSS) modulation. We restrict our attention to the 11 channels in use in North America.

To avoid interference, wireless radios are expected to operate on non-overlapping channels, i.e., channels separated by at least 22 MHz. For example, if two APs are operating on the same channel in a wireless cell, then their signals will interfere with each other. The same applies to any other radiating device, such as a microwave oven or cordless phone. From Figure 2, we observe that the following channel combinations do not overlap with each other: $\{1,6,11\}$, $\{2,7\}$, $\{3,8\}$, $\{4,9\}$, and $\{5,10\}$. Channels 1, 6, and 11 are the most commonly used non-overlapping channels in Wi-Fi deployments.




Fig. 4. Duty cycle analysis for six non-Wi-Fi devices

3.2 Metrics

We use spectrograms and duty cycles to characterize interferers at the physical layer. The spectrogram is a representation of the RF power levels over time in the spectrum. Each vertical line in the spectrogram shows the RF power as a function of frequency measured over a time interval of 1 second. Spectrograms offer a temporal perspective of RF power in the frequency domain. The duty cycle measures the RF power in the spectrum. In this work, duty cycle is calculated by measuring the percentage of time the RF signal is 20 dBm above the noise floor. Duty cycle is an indicator of the impact of RF power on network performance. We next describe the physical-layer characteristic of each interferer in isolation.

3.3 Measurement Results

Figure 3 shows the spectrograms for the interferers. The X-axis represents the time period of the measurements. The Y-axis tic marks represent the centre frequencies of the even numbered Wi-Fi channels; i.e., channels 2, 4, 6, 8, 10, 12, and 14. The colour contour lines represent the power levels of the signal, where red indicates the strongest and blue the weakest power levels. When interpreting these graphs, it is important to note that each device may use a different range of RF power levels and the same colour therefore may refer to a different power level in the different graphs. Figure 4 shows the duty cycle FFT measurements for the interferers. The bottom X-axis tic marks represent the centre frequencies of the even numbered Wi-Fi channels, while the top X-axis tic marks show the channel numbers corresponding to those centre frequencies.

Microwave Oven: Figure 3(a) shows that the microwave oven affects about half the available channels (6-12) in the 2.4 GHz band with the highest energy concentrated on channel 9. Note that the microwave oven is operating at high power levels of -80 to -60 dBm. RF power levels above -80 dBm are enough to cause interference with Wi-Fi networks. Figure 4(a) shows that the average duty cycle of the microwave is 50%. Since the microwave signal sweeps through a wideband of the spectrum and has a high average duty cycle, it is likely to affect nearby Wi-Fi devices. Microwave ovens are created with a shield such that all radiation is restricted to the oven cavity, however, with use over time these ovens can leak some radiation.

Analog Wireless Video Camera: Figure 3(b) illustrates the narrowband continuous transmitting nature of an analog video camera. The transmit power is similar to that of the microwave oven, however, in this case this energy is concentrated on a very small portion of the spectrum (channels 4-8). Figure 4(b) shows that the duty cycle of the analog video camera reaches 100% indicating that no Wi-Fi device in the vicinity will be able to operate on channels 4-8. Because of its continuous transmission nature, this device can cause prolonged periods of service disruption.

Analog Cordless Phone: The analog phone concentrates most of its energy on channel 3 in Figure 3(c). The figure also highlights the narrowband fixed-frequency transmission nature of the analog phone. Figure 4(c) shows that the analog phone has duty cycle as high as 85%. The high duty cycle of the analog phone indicates that it will severely impact Wi-Fi operation. Analog phones are quickly becoming out-dated, however, they are still available for purchase and are widely used around the world.

Digital Cordless Phone: The frequency hopping feature of the digital phone is illustrated in Figure 3(d). This phone utilizes DSS, a newer technology than the one used by analog phones. The phone continuously changes channels, only staying on a portion of the spectrum for a small amount of time, reducing its interference. In Figure 4(d) we observe that the maximum duty cycle is 4.5%, indicating that this device may not severely interfere with Wi-Fi devices.⁵

Bluetooth Headset: The Frequency Hopping Spread Spectrum (FHSS) nature of Bluetooth is highlighted in Figure 3(e). The Bluetooth device hops across all the channels. Although the energy emitted by the Bluetooth device may appear high, its duty

⁵ More recently, Digital Enhanced Cordless Telecommunications (DECT) phones have become popular. These phones operate on the 1.9 and 5.8 GHz band and do not interfere with 802.11 devices.

cycle values are much lower. Furthermore, FHSS technology is limited to 2 Mbps and does not consume much of the available bandwidth. Figure 4(e) shows that the maximum duty cycle attained is 3.5%, which may not affect Wi-Fi devices seriously.

Wireless Jammer: Figure 3(f) shows the wideband characteristics of a wireless jammer. The jammer emits signals on all channels at high power levels in quick succession. Figure 4(f) shows the duty cycle varies between 10% and 60%. Wi-Fi devices use Clear Channel Assessment to sense when the channel is clear for transmission. The wideband jammer ensures that the RF medium is never clear, thus preventing Wi-Fi devices from functioning properly.

4 Impact of Interferers on Wi-Fi Traffic

In this section, we study the impact of interferers on Wi-Fi traffic. In particular, we study the impact of interferers on data traffic, video traffic, and voice traffic. The first experiment for each workload was conducted in an interference-neutral environment. Next, we repeated the experiments where the Wi-Fi link was subjected to interference from one specific interferer at a certain distance. The distances used were 5, 10, 15, 20, 25, and 30 m. We did not consider the digital phone since its physical layer characteristics are similar to that of Bluetooth. Each experiment was performed four times, and we report the average percentage difference between the baseline performance (no interference) and that when the experiment was subjected to interference. For the data experiment, we used the throughput Quality of Service (QoS) metric. For the video and voice experiments, we used a Quality of Experience⁶ (QoE) metric called Mean Opinion Score (MOS).

4.1 Experimental Workload Traffic

Data Traffic: We used the Iperf⁷ tool to measure throughput of the Wi-Fi link. We ran Iperf in server mode at the server and in client mode at the four wireless client laptops. Iperf can run throughput tests using TCP or UDP packets. We used the TCP option for our experiment. For the TCP tests, Iperf requires the user to set an appropriate TCP window size. If the window size is set too low, the throughput measurements may be incorrect. We found that a window size of 148 KB was sufficient to properly measure the throughput of the Wi-Fi link. The workload consisted of creating bidirectional TCP traffic between the server and the four clients for a period of 3 minutes.

Video Traffic: Video (or Voice) quality can be quantified using subjective methods such as MOS. For our experiments, we used a hybrid subjective assessment scheme called Pseudo Subjective Quality Assessment (PSQA) [7, 8]. PSQA uses random neural networks to automatically calculate MOS of a video or voice sample. We set up the server workstation to stream a 3-minute video using the VLC⁸ media player. The (normal quality standard-definition) video had a resolution of 624×352 , frame rate of 29 frames per second, and was encoded using the Xvid⁹ codec at a bit rate of 656 Kbps.

⁶ http://en.wikipedia.org/wiki/Quality_of_experience

⁷ <http://iperf.sourceforge.net/>

⁸ <http://www.videolan.org/vlc/>

⁹ <http://www.xvid.org/>

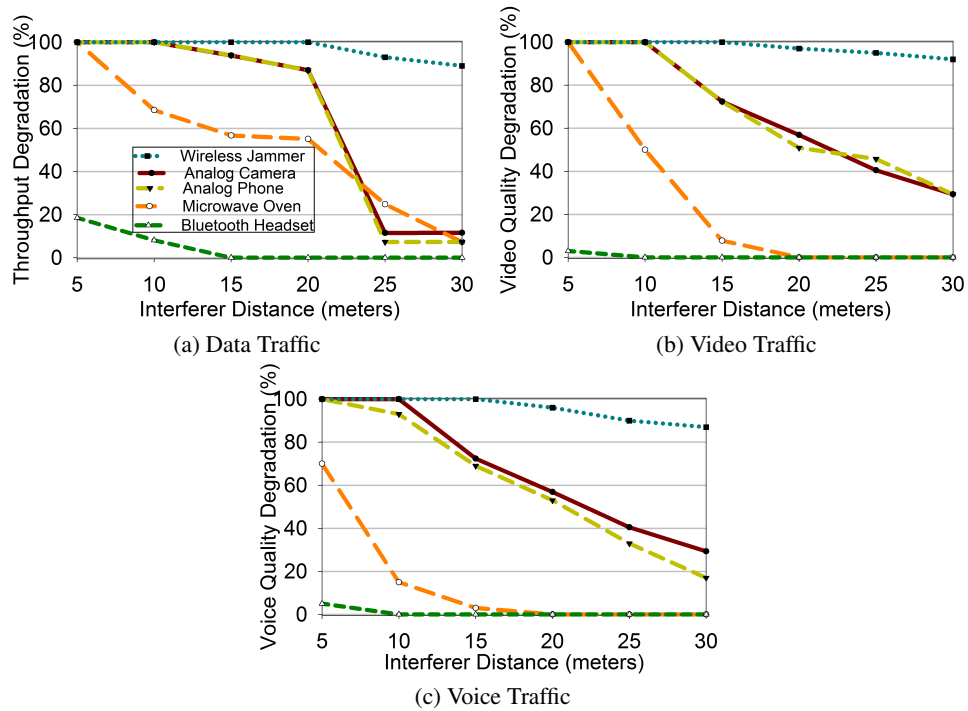


Fig. 5. Impact of interference on different workloads

The VLC players on the four clients were configured to receive this video stream. We collected the MOS measurements using PSQA from the client-side.

Voice Traffic: We recorded a 3-minute VoIP conversation between two people. We next separated the voices of the two people and created two different audio files. Recall our voice experiment setup had two pairs of communicating laptops. After establishing a direct connection via the Wi-Fi link for a pair of laptops, we played one audio file on one end and the second audio file on the other end. The two audio files were synchronized such that at no time the two audio files played a human voice simultaneously. The (telephone quality) audio files were encoded using the FLAC¹⁰ codec at a sampling rate of 16 kHz and a bit rate of 56 Kbps. We used a workstation connected to the AP to collect the MOS measurements using PSQA of the voice communication.

4.2 Experimental Results

Figures 5(a), (b) and (c) show the degradation in data throughput, video quality, and voice quality, respectively, in the presence of interference.

As the diverse nature of these workloads require different quality measures, direct comparisons of the degradation of the different services require some care. Voice traffic, which typically uses smaller packets, handles interference the best (with the metrics

¹⁰ <http://flac.sourceforge.net/>

used here). As expected, we consistently notice close to 100% degradation for our experiments using the wireless jammer. This is consistent up to 20 m, beyond which we observe a slight decline in its impact. In the following, we discuss the impact of the unintentional interferers on the different workloads.

Data Traffic: The Bluetooth headset reduced the throughput by 20% at close distances. The degradation may be surprising, since Bluetooth devices have low duty cycle and are designed to accommodate Wi-Fi devices; however, it is less serious compared to other interferers. For example, the microwave oven resulted in zero throughput at close distance. While its impact declined gradually as it was moved away from the Wi-Fi network, it caused a 25% degradation in throughput even at a distance of 25 m. The analog phone and video camera are both continuous transmitters. Hence, their impact on throughput is similar. Their interference is significant at close distance, however, there is a sharp decrease at distances beyond 20 m.

Video Traffic: Although Bluetooth had some impact on data traffic, there was minimal impact on video traffic. The microwave oven at close distance severely disrupted the video stream, while at longer distances it reduced the video quality by only 10%. The analog camera and analog phone had similar impacts on the video stream. Even at longer distances, they reduced the video quality by 50%.

Voice Traffic: Bluetooth had minor impact at short distance and no impact at longer distances. The microwave oven caused approximately 75% degradation at close range, and this plummeted as the distance increased. The analog phone and video camera had severe impact at close range, but this consistently decreased with longer distances. They still caused about 30% degradation at a distance of 30 m. The impact of interference is comparatively lower for voice traffic.

5 Interference in Campus Network

We used the Spectrum Analyzer to study the channel utilization of our campus network. We took physical-layer measurements for 8 hours during a weekday (Tuesday, January 20, 2009). The passive measurements were taken at a popular location in the campus frequented by students and faculty members. We report our results using the channel utilization metric. Channel utilization is the percentage of time a transmission is present from a known RF source in a given channel. It helps us understand how much of the channel is available for use at a given time.

Figure 6(a) shows the channel utilization of the three channels used in our campus network. The figures record the channel utilization by all devices (Wi-Fi and Interferers) transmitting in the 2.4 GHz ISM band.¹¹ We observe that channels 1 and 11 had heavier usage in comparison to channel 6. In case of channels 1 and 6, utilization peaked near 60%, while for channel 11 it was over 90%. After close inspection, we found that most of the spikes were caused by interferers.

On channel 11, the spikes are the most alarming. We focus on three spikes and try to understand the interferers involved. These spikes are observed at times 12:45, 15:38, and 17:17. We found that each of these spikes was mainly caused due to microwave

¹¹ The Spectrum Analyzer cannot separate the Wi-Fi and non-Wi-Fi channel utilization over time, but it did allow to us to investigate the utilization at particular points in time.




Fig. 6. Passive measurements from campus Wi-Fi network

ovens and cordless phones operating simultaneously in the area of the Wi-Fi network. For the first spike, we also noticed some other fixed-frequency interferers, whereas the second and third spike were highly influenced by Bluetooth communication. In total, the interferers consumed 80% of the channel during the first spike. The second and third spike used much less of the channel.

Figure 6(b) shows the average number of interferers observed on channel 11 as a function of time. We observed 100s of interference occurrences involving non-Wi-Fi devices on channel 11. These include 22 instances of Bluetooth device pairs communicating (6 unique piconets), 3 instances of analog cordless phones (base and handset), 71 instances of digital cordless phones (30 bidirectional base and handset instances and 41 cases where only the base station was observed), as well as 46 instances involving other unclassified interferers.

Overall, we found that microwave ovens and cordless phones were surprisingly active throughout the trace period, while Bluetooth piconet instances were temporally clustered. Furthermore, we note that most active times for the various interferers were relatively short, but with a few long instances. For example, the average active period for microwave ovens (2:12), Bluetooth (3:05), digital cordless phones (7:32) and other interferers (0:42) were relatively short compared to the corresponding devices' longest observed active periods (16:37, 22:27, 2:28:01, 5:02).¹²

Figure 6(c) shows the cumulative active time (i.e., the time the interferers were transmitting) of interferers over the trace period. Microwave ovens were active 31% of

¹² The active period length is reported in *hours:minutes:seconds*.

the total trace period and had an average (and maximum) duty cycle of 9% (40%). The average signal strength was approximately -80 dBm, which is enough to cause some interference. Bluetooth devices were active 12.5% of the time and had an average (low) duty cycle of roughly 5%. With digital base stations always transmitting (at a duty cycle $< 7\%$) when not connected with the handset, it was not surprising that we found that cordless phones were cumulatively active for more than the entire trace period. We note, however, that the average (maximum) duty cycle increased to 15% (20%) when the base and handset were connected. Other interferers were observed active only 6% of the total trace period, but had the highest average duty cycle (15%), maximum duty cycle (60%), and signal strength (-43 dBm) of all the interferers observed in the campus Wi-Fi network.

The exact channel usage of these observed non-Wi-Fi devices may differ from that of the particular devices studied in our controlled experiments, however, the general characteristics are similar. With a small number of heterogeneous non-Wi-Fi devices active at each point in time, it is not surprising that we do not find strong correlation between the average number of active devices and the peaks observed above (e.g., comparing Figures 6(a) and 6(b)). Instead, individual device characteristics and the distance of devices from the network are likely to have the most impact. Referring back to the impact of device and distance (Section 4), it is clear that a few non-Wi-Fi devices can have a huge effect on the performance of the Wi-Fi traffic, especially if located close to the network. Furthermore, while our measurements were taken (in a student/faculty lounge) away from the campus food court and office areas, the interference from microwave ovens and cordless phones would likely be even more significant in such areas. We leave comparisons of the non-Wi-Fi interference in different areas as future work.

6 Related Work

Interference in wireless networks has been investigated mostly in the context of co-existing Wi-Fi networks, Bluetooth networks, and microwave ovens.

Taher *et al.* [14] used laboratory measurements to develop an analytical model of microwave oven signals. Karhima *et al.* [12] performed measurements on an ad-hoc wireless LAN under narrowband and wideband jamming. They found that in case of wideband jamming, 802.11g can offer higher transmission rates than 802.11b, when the packet error rate is the same. They also found that 802.11g was more prone to complete jamming, while 802.11b could still operate at lower transmission rates due to its DSSS modulation scheme. Golmie *et al.* [9] explored the mutual impact of interference on a closed loop environment consisting of Wi-Fi and Bluetooth networks. They found that even by sufficiently increasing the transmission power levels of the Wi-Fi network to that of the Bluetooth network could not reduce packet loss. Our work complements these studies. We use measurements to characterize a wide range of common non-Wi-Fi devices (including microwave ovens and Bluetooth) at the physical layer and use experiments to quantify their impact on different traffic workloads.

Vogeler *et al.* [15] presented methods for detection and suppression of interference due to Bluetooth in 802.11g networks. Ho *et al.* [11] studied the performance impact on a Wi-Fi network due to Bluetooth and HomeRF devices using simulations and described

design challenges for deployment of 5 GHz wireless networks. Our work is orthogonal to these papers. We focus on passive measurements to study interference from common non-Wi-Fi devices.

Gummadi *et al.* [10] studied the impact of malicious and unintentional interferers on 802.11 networks. They found that some of these interferers could cause considerable performance degradation for commercial 802.11 NICs. Using an SINR model, they noted that changing 802.11 parameters was not helpful, and proposed a rapid channel hopping scheme that improved interference tolerance for these NICs. Farpoint Group studied the effect of interference on general, video, and voice traffic in a Wi-Fi network from two vantage points: short-range (25 feet) and long-range (50 feet) [3–5]. In contrast, we present a comprehensive measurement study of impact of interference from non-Wi-Fi devices. We also used QoS and QoE metrics to quantify the degradation of video and voice quality in a Wi-Fi network due to interference. We observed that operational Wi-Fi networks may be subject to ambient interference effects at anytime.

7 Concluding Remarks

In this paper, we characterized the RF behaviour of non-Wi-Fi devices, analyzed the impact of interference on data, video, and voice traffic, and examined interference in a live campus network. Overall, we found that the campus network is exposed to a large variety of non-Wi-Fi devices, and that these devices can have a significant impact on the interference level in the network. Our controlled experiments showed that, even at distances up to 30 m, some of these non-Wi-Fi devices can have a significant negative impact on data, video, and voice traffic. While microwave ovens, wireless analog video cameras, and analog cordless phones typically have the most significant negative impact on Wi-Fi networks, the performance degradation due to a digital cordless phone or a Bluetooth device (which tries to be Wi-Fi friendly) can be noticeable (e.g., 20% at close distances). With the campus network (and likely many other home and enterprise networks) being highly exposed to different types of unintentional interferers, it is important to find (new) ways to identify and mitigate non-Wi-Fi interference.

Network practitioners often use link-layer and transport-layer statistics to investigate interference in a Wi-Fi network. We note that these alone should not be the means for troubleshooting a network. We observe that physical-layer characteristics may be used as primary indicators to identify and mitigate interferers [2]. In general, we believe that interference can be mitigated by identifying and removing the interfering device (if possible) or shielding interferers instead. Careful channel selection at APs may be helpful (e.g., in both our controlled and real world measurements we found that microwave ovens primarily affected channels 6 and higher). APs may be fitted with sensors to detect interference and switch channels automatically [1]. Using multi-sector antennas [13] and controlling data rates to avoid false backoffs can make the network more interference-resilient (although this is a trade-off since the lower data rates allow more noise-immune communication). We conclude by noting that there are many different non-Wi-Fi devices that may cause interference in Wi-Fi networks, and it is therefore important to understand and quickly adapt to the devices affecting the performance of the Wi-Fi channels.

8 Acknowledgements

Financial support for this research was provided by Informatics Circle of Research Excellence (iCORE) in the Province of Alberta, as well as by Canada's Natural Sciences and Engineering Research Council (NSERC). The authors also thank the anonymous reviewers for their constructive comments on an earlier version of this paper.

References

1. Aruba Networks: Advanced RF Management for Wireless Grids. Technical Brief (2007), Available online at <http://tinyurl.com/aruba-report>
2. Cisco: 20 Myths of Wi-Fi Interference: Dispel Myths to Gain High-Performing and Reliable Wireless. White Paper (2007), Available online at <http://tinyurl.com/cisco-report>.
3. Farpoint Group: The Effects of Interference on General WLAN Traffic. Technical Note FPG 2006-328.3 (January, 2008), Available online at <http://tinyurl.com/FarpointGeneral>
4. Farpoint Group: The Effects of Interference on Video Over Wi-Fi. Technical Note FPG 2006-329.3 (January, 2008), Available online at <http://tinyurl.com/FarpointVideo>
5. Farpoint Group: The Effects of Interference on VoFi Traffic. Technical Note FPG 2006-330.3 (January, 2008), Available online at <http://tinyurl.com/FarpointVoFi>
6. Flickenger, R.: Building Wireless Community Networks. O'Reilly (2003)
7. Ghareeb, M., Viho, C.: Performance Evaluations of a QoE-Based Multipath Video Streaming Mechanism over Video Distribution Network. In: Proc. of International Workshop on Future of Multimedia Networking (FMN). pp. 236–241. Coimbra, Portugal (June, 2009)
8. Ghareeb, M., Viho, C., Ksentini, A.: An Adaptive Mechanism For Multipath Video Streaming Over Video Distribution Network. In: Proc. of International Conference on Advances in Multimedia (MMEDIA). pp. 6–11. Colmar, France (July, 2009)
9. Golmie, N., VanDyck, R., Soltanian, A., Tonnerre, A., Rebala, O.: Interference evaluation of Bluetooth and IEEE 802.11b Systems. *Wireless Networks* 9(3), 201–211 (May, 2003)
10. Gummadi, R., Wetherall, D., Greenstein, B., Seshan, S.: Understanding and Mitigating the Impact of RF Interference on 802.11 Networks. In: Proc. of ACM SIGCOMM International Conference. pp. 385–396. Koyoto, Japan (August, 2007)
11. Ho, M.J., Rawles, M., Vrijkorte, M., Fei, L.: RF Challenges for 2.4 and 5 GHz WLAN Deployment and Design. In: Proc. of IEEE Wireless Communications and Networking Conference (WCNC). pp. 783–788. Orlando, U.S.A. (March, 2002)
12. Karhima, T., Silvennoinen, A., Hall, M., Haggman, S.: IEEE 802.11 b/g WLAN Tolerance to Jamming. In: Proc. of Military Communications Conference (MILCOM). pp. 1364–1370. Monterey, U.S.A. (October, 2004)
13. Lundgren, H., Subramanian, A., Salonidis, T., Carrera, M., Guyadec, P.L.: Interference Mitigation in WiFi Networks using Multi-sector Antennas. In: Proc. of ACM international Workshop on Experimental Evaluation and Characterization (WiNTECH). pp. 87–88. Beijing, China (September, 2009)
14. Taher, T., Al-Banna, A., Ucci, D., LoCicero, J.: Characterization of an Unintentional Wi-Fi Interference Device - The Residential Microwave Oven. In: Proc. of Military Communications Conference (MILCOM). Washinton D.C., U.S.A. (October, 2006)
15. Vogeler, S., Brotje, L., Kammeyer, K.D., Ruckriem, R., Fechtel, S.: Suppression of Bluetooth Interference on OFDM in the 2.4 GHz ISM Band. In: Proc. of International OFDM-Workshop (InOWo). Hamburg, Germany (September, 2003)