Experimental Verification of a Radio-Frequency Power Model for Wi-Fi Technology

Minyu Fang and David Malone∗

∗ Hamilton Institute, NUI Maynooth, Kildare, Ireland. Fax: +353-1-708 6296. Tel: +353-1-708 6463. Email: David.Malone@nuim.ie. Work supported by Science Foundation Ireland Grant RFP-07-ENEF530.

This is a non-final version of the article published in final form in Health Physics. Vol 98, Issue 4, 574-583, April 2010. See http://www.health-physics.com/

Abstract

When assessing the power emitted from a Wi-Fi network, it has been observed that these networks operate at a relatively low duty cycle. In this paper, we extend a recently introduced model of emitted power in Wi-Fi networks to cover conditions where devices do not always have packets to transmit. We present experimental results to validate the original model and its extension by developing approximate, but practical, testbed measurement techniques. The accuracy of the models is confirmed, with small relative errors: less than 5–10%. Moreover, we confirm that the greatest power is emitted when the network is saturated with traffic. Using this we give a simple technique to quickly estimate power output based traffic levels and give examples showing how this might be used in practice to predict current or future power output from a Wi-Fi network.

Keywords

exposure, radiofrequency; dose assessment; safety standards; computers

Introduction

Wireless Local Area Networks (WLANs) have become ubiquitous in recent years. These networks use radio frequency (RF) energy and regulations stipulating the maximum transmit power used by WLANs are set by the FCC (Federal Communications Commission) and CEPT (European Conference of Postal and Telecommunications Administrations). Signals are transmitted at low powers, typically 0.1 W for both computers and access points (APs). The acceptable thresholds for absorbed radiated power given by the ICNIRP (International Commission on Non-Ionizing Radiation Protection) are 80mWkg^{-1} for the general public and 400mW kg^{-1} for occupational exposure (whole body).

Results so far show that exposures from standard deployments are well within internationally accepted ICNIRP guidelines (Schmid et al. 2007 and Foster 2007). While

Wi-Fi has not attracted the same level of interest as mobile phone networks, there still exists public concern regarding health and safety issues, particularly in schools (Peyman 2009) but also in homes and offices (ACRBR 2008, Kühn 2006). The chairman of the UK's Health Protection Agency (a body established to protect the public from environmental hazards, including non-ionizing radiation) has stated it would be timely to carry out further research as this technology is rolled out. The trend is toward denser Wi-Fi deployments, such as extremely dense hotspots in an urban area. A house or apartment could have ten Wi-Fi devices, including broadband routers, laptops, phones, PDAs, games consoles and media players. Classrooms or conference halls could have larger numbers of devices with higher levels of activity. It is therefore useful to model and evaluate how the radiated power scales with the number of stations and level of activity, to determine if radiation levels are within acceptable limits. Such models may be of use for both retrospectively assessing RF levels or for planning of future WLAN use, where measurement is not possible.

In this paper we aim to give some models for estimating the power output of a wireless LAN that is not always busy. As noted by various authors (e.g. ICNIRP 2009), Wi-Fi transmissions are intermittent and time-averaged powers depend on the amount of data transferred; it is this issue we consider. Factors such as the speed of broadband access links and the speed at which people can navigate the network serve to restrict how busy a WLAN can become. For example, an architect's office might send large files to clients each day, but be restricted by a broadband link. Alternatively, someone watching YouTube videos will tend not to download faster than they can watch them. Of course, the wireless link may become very busy in cases with fast links (e.g. a large school/campus) or where network transfers are local (e.g. backing data up to a local server). We will look at both cases of non-saturated and saturated networks.

Most deployments of WLANs are based on widely used Wi-Fi technology. In this paper, we concentrate on WLANs based on the IEEE 802.11 standard and its amendments. We do not consider other technologies such as WiMAX (Worldwide Interoperability for Microwave Access) or 3G. In the typical Wi-Fi scenario (known as an infrastructure mode network), every station communicates with an access point (AP) connected to the wired Internet network. Widely deployed IEEE 802.11b and IEEE 802.11g operate in the unlicensed spectrum at 2.4 GHz; IEEE 802.11a utilizes spectrum around 5 GHz. We focus on 802.11b in the 2.4G Hz band, because it is a least common denominator and supported by almost all existing hardware. Nonetheless, our results extend directly to 802.11a and 802.11g.

Wi-Fi devices based on the IEEE 802.11 standard only transmit and radiate power when they have a data packet to send and when they are permitted to do so by the 802.11 protocol. The 802.11 Medium Access Control (MAC) protocol regulates channel access. It is the impact of this MAC on transmitted power that will be of interest to us. If more than one device transmits at a time the result is a collision, which results in no data being successfully transferred. Thus, the MAC attempts to control transmissions so that there is a high likelihood only one device transmits at a time. This is achieved by the MAC by inserting random time gaps (called backoff periods) after transmissions and collisions. Hence, the MAC protocol has an important impact on transmitted power by 802.11 devices, achieving a middle ground between all devices transmitting at once and just one device transmitting at a time. An 802.11 MAC model (Bianchi 2000) was developed to determine the performance, where the transmission probability and collision probability can be calculated as a function of the number of stations, assuming that each device always has a packet to send. Based on analysis of these probabilities, Malone and Malone 2009 estimated the mean transmitted power as a function of the number of stations under the same assumptions and studied the total power emitted in error-free, error-prone, broadcast and unicast networks.

To consider unsaturated networks, we extend the model of power output beyond the saturated situation using the non-saturated models from Duffy and Ganesh 2007 and Malone et al. 2007. These models allow the amount of network traffic at each device to be varied. The two models consider two extremes: Duffy and Ganesh 2007 assume that traffic arriving while the device is busy will be queued until it can be transmitted, whereas Malone et al. 2007 assume such traffic is discarded without being transmitted. We call Duffy and Ganesh 2007 the infinite buffer model and Malone et al. 2007 the no-buffer model. We also present experimental results to compare the model predictions for the saturated model presented in Malone and Malone 2009 and the non-saturated models that are described in this paper. We will apply these models to consider the power output associated with a number of scenarios.

Note that we calculate the sum of the power of all stations in the network, rather than the exposure at a particular point. Like Malone and Malone 2009 we omit some important factors for calculating exposure, such as the distances between devices, reception errors caused by absorption/reflection in the environment or interference from other devices sharing the same frequency and so on. Similarly, we assume that maximum transmitter power approximates actual transmit power. These assumptions provide upper bounds of the transmitted power involved in exposure and are considered further in our discussion.

The rest of this paper is organized as follows. We first outline the testbed system employed in our experiment and our experimental technique in the method section. A summary of the theoretical analysis of the transmitted power model is presented in the section on modeling power, including the extension to non-saturated conditions. Then, the experiment results are provided and compared to the theoretical counterparts in the results section. A discussion of these results, along with quick techniques for estimating output power then follows.

Method

Experiments are carried out on our wireless testbed as shown in Fig. 1, which is configured in the infrastructure mode. This is a similar configuration to that which might be found with a number of devices in a home, or in a public hot spot. Each station is representative of a laptop or other wireless device. The testbed includes 9 stations that are a collection of PC-based embedded Linux boxes based on Soekris net4801 (Soekris 2004), one desktop PC acting as a client station, and another desktop PC acting as an access point. All devices are installed with a standard wireless card (an Atheros AR5215 802.11b/g PCI card) and an external antenna. All stations, including the AP, run a Linux 2.6.8.1 kernel and a version of the MADWIFI wireless driver, which we have modified to allow greater logging and control (MADWIFI 2009). Meanwhile, these stations are connected through 100Mbps wired Ethernet to a PC that controls the testbed system. The desktop PC is employed as a station to record detailed per-packet statistics. An advantage of this PC acting as a station is that there is adequate storage space, and sufficient RAM and CPU for the collection of statistics. We use a number of tools common in the traffic engineering community in our testbed. The **mgen** tool is used to generate UDP traffic (NRL 2009). We use the Linux **sysctl** command to specify traffic parameters, such as fixed data rate of the wireless card. The Linux **ssh** and **scp** commands are used for the network management and the control of traffic sources over wired Ethernet ports. We also use the **athstats** tool to collect statistics from the wireless driver (MADWIFI 2009).

There were two measurement methods employed in the experiments. In the first method we recorded the number of successful transmissions and collisions by analyzing a trace file produced by the modified driver and stored in the desktop PC station. Then, we scaled up the number of transmissions/collisions by the number of stations in a wireless network to approximate the results of the whole network. Here we assumed that the network is symmetric and so other stations would have the same performance as the desktop PC. This technique only required us to record data at one PC, but we only expected good accuracy if the traffic load on the stations was symmetric and they were in a symmetric environment. We will see the implications of this in our results.

Our second measurement technique used **athstats**, which recorded basic statistics relating to the wireless card. We focused on the number of transmitted frames, the number of retries and the number of failed transmissions. We recorded these statistics for each station in our testbed. Compared to the first method, we expected higher accuracy, as we have a picture of the whole system's performance. This method does not require the network is to be symmetric. In the following sections we will show results generated by both of the methods.

In our tests, we configured all stations identically to make the network symmetric. Regardless, there still existed some differences due to the environment. An example, depicted in Fig. 2 and Fig. 3, shows the total number of retries and transmissions as the load was varied. Results are shown for 9 stations operating at the same time. It is evident that most of the stations are relatively similar in terms of the numbers of the transmissions and retries. One station showed a much smaller number of retries. We will see in the results section that this asymmetry actually has a small impact in the prediction of transmitted power (which, in this case, is dominated by the number of successful transmissions rather than the number of retries).

After obtaining the desired statistics, the transmitted power was calculated as

$$\frac{n_s E_s + n_c E_c}{T_{total}} \tag{1}$$

where E_s and E_c are the mean energy associated with a successful transmission and a collision. These were calculated in the same way as for their theoretical counterparts, which will be described in the next section. n_s is the number of successful transmissions and n_c was the number of collisions. These were calculated from our experiments, as described above. T_{total} is the time for the whole experiment, which was calculated by subtracting the first in-queue time from the last in-queue time.

Model of Transmitted Power

The IEEE 802.11 MAC defines two different access mechanisms, the mandatory Distributed Coordination Function (DCF) and Point Coordination Function (PCF). PCF provides centrally controlled channel access through polling, but is rarely used in practice. The models we are interested in are of DCF, so we briefly explain it here. For a more complete and detailed presentation, refer to the 802.11 standard (IEEE 1997).

DCF is based on Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). It works as a listen-before-talk protocol. On sensing the transmission medium idle more than the DCF Inter-frame Space (DIFS), a station may start to transmit or, if in backoff, it may count down. When any station transmits, other stations wait until the medium becomes idle again at least for DIFS. When a destination successfully receives a frame, it will acknowledge by sending an ACK (Acknowledgement) frame after SIFS (Short Inter-frame Space). If two or more transmissions start at the same time, the result is a collision, which typically results in no packet being successfully received. Packets can also be lost due to noise, fading or other radio-frequency effects.

The backoff procedure involves counting down a randomly chosen number of slots where the medium is idle (usually each 20us long). The number of slots is chosen uniformly in $\{0, 1, \ldots CW - 1\}$, where CW is the contention window and depends on the number of retransmissions. The initial CW is set at CW_{min}, typically 32. The value of CW doubles on a collision up to a maximum value of CW_{max}, typically 1024. After a successful transmission CW is reset to CW_{min}. In our equations, we denote $W = CW_{min}$ and $2^m W = CW_{max}$.

In Bianchi 2000 a mean field Markov Chain model was established to obtain the performance of 802.11 DCF as a function of the number of station in saturated conditions. Based on the analysis of DCF's performance, the transmission probability was calculated and the mean power was obtained by Malone and Malone 2009 as

$$P = \frac{E}{T} = \frac{(E_S - E_C)n\tau(1-\tau)^{n-1} + E_C n\tau}{T_I(1-\tau)^n + T_S n\tau(1-\tau)^{n-1} + T_C(1-(1-\tau)^n - n\tau(1-\tau)^{n-1})} \quad (2)$$

where τ denotes the transmission probability, T represents the average length of a slot time, T_I is the length of an idle state, n is the number of stations in the network, T_S is the mean time for a successful transmission and T_C is the mean time for the length of a transmission. They are easily calculated from the 802.11 standards and network settings. At the same time, E_S and E_C can be estimated by introducing the nominal power output P_0 (say 100mW) and the length of time spent on successful transmission, T_{Es}, and collision transmission, T_{Ec}. They are given as below

$$E_S = \mathrm{P}_0 T_{Es} = P_0(2 * preamble + (header + payload)/rate + ack) \quad (3)$$

$$E_C = \mathrm{P}_0 T_{Ec} = P_0(preamble + (header + payload)/rate) \quad (4)$$

where $preamble$, $header$, $payload$ and ack are the times/sizes used for each of these transmissions and $rate$ is the speed at which data is transmitted. The calculation of these quantities is described in detail in the appendix of Malone and Malone 2009, and depends on packet lengths, protocol constants and so on. Quantities such as the signal propagation delay also have a small impact. The values used in this paper are shown in Table 1.

We can also give an expression for the duty cycle for RF energy from Wi-Fi devices based on IEEE 802.11. The duty cycle for the network, which is the fraction of the time during which at least one station is transmitting is given by:

$$D_{net} = \frac{T_{Es} n\tau(1-\tau)^{n-1} + T_{Ec}(1-(1-\tau)^n - n\tau(1-\tau)^{n-1})}{T_I(1-\tau)^n + T_s n\tau(1-\tau)^{n-1} + T_c(1-(1-\tau)^n - n\tau(1-\tau)^{n-1})} \quad (5)$$

where T_{Es} and T_{Ec} respectively present the effective time of a successful transmission and a collision when a station is transmitting. The duty cycle for a single station's activity will be

$$D_{sta} = \frac{\tau(1-p)T_{Ec} + \tau p T_{Ec}}{T_I(1-\tau)^n + T_S n\tau(1-\tau)^{n-1} + T_C(1-(1-\tau)^n - n\tau(1-\tau)^{n-1})} \qquad (6)$$

We note that the energy associated with a given station can be obtained by multiplying the duty cycle by the nominal power. The power for the network can then be obtained by summing the power outputs over all stations, to give the same result as equation (2).

All the models that we look at assume that there is a fixed collision probability p. For simplicity, we restrict our attention to the case where the network is symmetric, so all stations have the same p. The number of devices is n. Note that $(1-\tau)^{n-1}$ represents the probability that $n-1$ stations do not transmit. Thus all our models use the relationship

$$(1-\tau)^{n-1} = 1-p. \qquad (7)$$

Saturation Model

In the saturation model of Bianchi 2000, the transmission probability τ is calculated as

$$\tau = \frac{2(1-2p)}{(1-2p)(W+1) + pW(1-(2p)^m)} \qquad (8)$$

This is the main expression used in Malone and Malone 2009 to calculate τ.

Non-Saturation Model

The non-saturated model Malone et al. 2007 is based on an idealized assumption of no buffering. This is achieved by assuming q, the probability of a packet arriving at the MAC during an average slot time, equals to r, the probability that at least one packet arrived while the station is idle. The expression τ is calculated in terms of a normalization factor $b_{(0,0)e}$,

$$\begin{aligned}
\frac{1}{b_{(0,0)e}} &= (1-q) + \frac{q^2 W(W+1)}{2(1-(1-q)^W)} \\
&+ \frac{q(W+1)}{2(1-q)}\left(\frac{q^2 W}{1-(1-q)^W} + p(1-q) - q(1-p)^2\right) \\
&+ \frac{pq^2}{2(1-q)(1-p)}\left(\frac{W}{1-(1-q)^W} - (1-p)^2\right)\left(2W\frac{1-p-p(2p)^{m-1}}{1-2p} + 1\right)
\end{aligned} \tag{9}$$

and then

$$\tau = b_{(0,0)e}\frac{q^2}{1-q}\left(\frac{W}{(1-p)(1-(1-q)^W)} - (1-p)\right) \tag{10}$$

An expression is given for *q,* the traffic arrival rate, in terms of λ, the packet arrival rate, using a Poisson traffic model (Bertsekas and Gallagher 1987):

$$q = 1 - e^{-\lambda T} \tag{11}$$

where T is the mean state time (the denominator in equation 2). Other traffic models are considered by Malone et al. 2007, but they are found to have similar performance in terms of throughput and collision probability. Hence, we only consider Poisson traffic here.

In addition to the model described above, we also consider another model with an infinite buffer introduced by Duffy and Ganesh 2007. The expression for τ becomes

$$\frac{1}{b_{(0,0)e}} = (1-q) + \frac{q^2 W(W+1)}{2(1-(1-q)^W)}$$
$$+ \frac{W+1}{2(1-r)}\left(\frac{q^2 rW}{1-(1-q)^W} + pq(1-r) + qr(1-p)^2\right)$$
$$+ \frac{p}{2(1-r)(1-p)}\left(\frac{q^2 W}{1-(1-q)^W} - rq(1-p)^2\right) \tag{12}$$
$$\left(2W\frac{1-p-p(2p)^{m-1}}{1-2p} + 1\right)$$

and

$$\tau = b_{(0,0)e}\frac{1}{1-r}\left(\frac{q^2 W}{(1-p)(1-(1-q)^W)} - rq(1-p)\right) \tag{13}$$

where r is obtained as

$$E(B(p)) = \frac{W}{2(1-2p)(1-p)}(1-p-p(2p)^m) \tag{14}$$

and

$$r = \min(1, -E(B(p))\log(1-q)) \tag{15}$$

where $E(B(p))$ represents the average MAC service time.

In summary, for saturated traffic, the relationship between the number of stations and the collision probability can be obtained by combining equation (7) and equation (8). From equation (2), we then calculate theoretical transmitted power for increasing numbers of stations. In the non-saturated no-buffer case, we use equation (9), equation (10) and equation (11) to find τ and then the theoretical transmitted power may be calculated using equation (2) for different numbers of stations and traffic loads. Similarly, the infinite-buffer case can be obtained with equation (12), equation (13), equation (14) and equation (15).

Results

The parameter values for our network are listed in Table 1. We use the values in Table 1, combined with analysis in the section where we model transmitted power, to compare the theoretical transmitted power with our measured results as we vary the number of stations and the offered load.

Fig. 4 shows a comparison between theoretical and experimental transmitted power in saturated conditions. In the experiments, UDP (User Datagram Protocol) is generated at the rate 11Mbps to saturate the network. The experiment is run for 100 seconds. As expected, power increases for larger numbers of stations. We see a good match between theory and experiment regardless of our measurement method: as predicted the power goes from slightly below the nominal value to around the nominal value as the number of stations is increased. Note that the results of experimental method 1 are slightly more variable. This is because the network is not symmetric in practice, as can been seen in Fig. 2 and Fig. 3. Using method 2, which more accurately reflects the total power actually transmitted, we see even better agreement with the model predictions.

Broadcast packets were also considered in Malone and Malone 2009, because the 802.11 backoff mechanism operates differently for packets that are destined to groups of devices. The differences arise from the fact that no ACK packet is sent, because no one station can know if the whole group has received the packet. We compare the predictions of the model with results in our test bed in Fig. 5. As expected, we see slightly higher power output than in Fig. 3, and the match between the theory and the testbed remains good.

The results in Fig. 4 and Fig. 5 show power being an approximately linear function of the number of saturated stations. The model has captured the intercept (80 mW) and slope (around 2 or 4 mW per station for non-broadcast/broadcast). For non-broadcast packets, we expect this slope to decrease for larger numbers of active stations, as the MAC's backoff will tend to reduce the transmission rate. This behavior is predicted by the model, as shown by Malone and Malone 2009, but our testbed is not large enough to verify the result.

In the following, we will focus on measurement method 2 because of its better accuracy. Fig. 6 shows the results of the big buffer experiments and Fig. 7 shows the results of the small buffer experiments. We look at the non-saturated case with an 11Mbps data rate and 100s experiment time. We approximate the infinite-buffer model with 200 packet buffer and the no-buffer model with a one packet buffer. We show the results of method 2 and equivalent model predictions for 2, 5 and 9 devices as we vary the load. The match between theory and experiment is excellent over lower traffic loads.

There is also good agreement as the network becomes saturated. For heavy load, the small buffer shows an almost perfect match for the cases of 2 stations and 5 stations, but underestimates by about 7% for 9 stations. By contrast, the big buffer is a better match to the theory line in the case of 9 stations, but slightly overestimates the power for 2 and 5 stations. In the intermediate region, larger discrepancies are possible.

We also present the results of our duty cycle calculations. Fig. 9 demonstrates the difference between the duty cycle of the entire network and the duty cycle summed over the stations as predicted by the model. Collisions allow the duty cycle summed over the stations to exceed 100%, which leads to the power exceeding the nominal value. Since our testbed results are per-station statistics, they only allow us to compare the duty cycle summed over stations with the model (Fig. 10 and 11). Fig.10 shows that the duty cycle saturated network increases quickly from roughly 75% at 1 station to roughly 105% at 9 stations. Fig.11 shows how the duty cycle is very small when the non-saturated 0.5 Mbps traffic is used, and increases linearly to around 6 Mbps. As these results are essentially rescaled versions of our power graphs, we see similarly good matches between model predictions and experimental results.

Discussion

Our experiments have confirmed that the models seem quite accurate. There are some differences observed between experiment and theory. However these might be explained by gaps between the assumptions of the theoretical model and the real world such as the network is not completely symmetric and infinite/no buffer were approximated with 200 packets or 1 packet. While the models are clearly not capturing the physical systems exactly, the predictions of radiated power would be accurate enough to make an informed dose calculation. Overall, the large buffer model's predictions appear more satisfactory for this purpose and are likely to better reflect the configuration of actual Wi-Fi devices.

When estimating the power output of a network, it may be useful to be able to estimate the largest possible power, regardless of traffic conditions. Intuition would suggest that the most power will be output when the network has the most to send, and is likely to be an implicit assumption of experimental studies. However, a feature of random-access MAC systems, such as 802.11, is that better data throughput can sometimes be achieved before the network becomes saturated. This is demonstrated, for example, in Malone et al 2007: for larger numbers of stations as load is increased the network's throughput increases to a peak and then decreases to its saturated level. Since throughput and power are loosely coupled, this might cast doubt on our Intuition. However, we see no power pre-saturation peak. We believe this is because, for realistic parameters, the expression for power (equation 2) will be an increasing function of transmission probability, unlike the expression for throughput. This suggests that, as we expect, the upper-limit on the outputted power is reasonably approximated by using the saturated network model.

Another useful observation from the graphs is that when there is a small amount of traffic in the network, the power is a linear function of the offered load. This is because the number of collisions for light loads is small, and so each packet is transmitted just once. Since 802.11 has a per-packet power overhead (for preamble, headers and ACK) and then a per-byte power cost (for transmitting the actual data) we may approximate the power as:

$$P = P_0\,(pps(2 * preamble + header/rate + ack) + bps/rate) \qquad (16)$$

where pps is the number of packets per second and bps is the number of bits per second. Fig. 8 shows the results of applying this rule of thumb to our experimental data. Note that the predictions are independent of the number of stations, and actually match well until the network reaches saturation.

Combining these two observations gives a simpler technique for estimating the power, if we know the amount of traffic. First, we use equation (16) to predict the power. Then we compare this to the power for a saturated network, and take the minimum. As examples, consider the following situations (details of the estimates are provided in the appendix).

1. An architect who uploads a large amount of data through their 1Mbps broadband link is concerned about their RF exposure. Using this technique we estimate that the power during the upload is about 17mW.
2. An office worker registers a complaint about a colleague who spends their lunch breaks watching YouTube videos, and is worried about the impact of the continuous downloads. By establishing the average rate associated with YouTube as 400Kbps, we are able to estimate a power of approximately 7mW.
3. In a high-school class, 30 students are encouraged to watch a short documentary from YouTube on their laptops at the end of each class. Parents express concern about 30 wireless devices being used at the same time. Using this technique, the network turns out to be saturated and our power estimate is about 120mW, rather than the potential 30*100mW.

Unsurprisingly, these powers are low when compared to the ICNIRP limit of 80mW kg^{-1}. However we now have a quick way to estimate power, given some information about the traffic in the Wi-Fi network.

In the paper, we have focused on the total power output. When more information was available, for example mixed output powers, distances from devices, antenna details, reflection patterns, etc, these could be incorporated into the model via the per-node duty cycle in equation (6). By weighting this duty cycle with per-node factors, such as output power, antenna gains and power decline due to distance, more exact dose calculations could be performed. Forster 2007 provides a detailed discussion of factors that could be accounted for.

In conclusion, we have extended the power model described by Malone and Malone 2009 to unsaturated networks. Through our testbed experiments, we have also verified this model and see a close match between theoretical predictions and experimental result. We find that the power of a saturated network is a reasonable upper bound on the power of an unsaturated network. For lightly loaded networks, we also offer a simple but accurate technique for approximating the power output. Finally, we give some examples of how these techniques might be applied.

References

Australian Centre for RF Bioeffects Research. EME in the Home. SW2008 Wireless and Health Unplugged and Uncertain. 2008.

Bertsekas, D.P. and Gallagher R.G. Data Networks. 1st ed. Longman Higher Education; 1987.

Bianchi G. Performance Analysis of IEEE 802.11 Distributed Coordination Function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, March 2000.

Duffy K and Ganesh AJ. Modeling the Impact of Buffering on 802.11. *IEEE Communications Letters*, 11(2):219–221, February 2007. Updated version available at http://www.hamilton.ie/ken_duffy/Downloads/commsl_kd_ajg.pdf.

Foster, K. R. Radiofrequency Exposure from Wireless LANs utilizing Wi-Fi Technology. *Health Physics*, 92(3):280–289; 2007.

Gill P, Arlitt M, Li S, Mahanti A. YouTube Traffic Characterization: A View From the Edge. ACM SIGCOMM Internet Measurement Conference, IMC'07.

IEEE. *Wirless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE std 802.11-1997 edition, 1997.

International Commission on Non-Ionizing Radiation Protection. Guidelines for Limiting Exposure to Time-varying Electric, Magnetic and Electromagnetic Fields. *Health Physics*, 74(4):494–522, April 1998.

International Commission on Non-Ionizing Radiation Protection. Exposure to High Frequency Electromagnetic Fields, Biological Effects and Health Consequences (100 kHz-300 GHz). ICNIRP 16/2009. Available at http://www.icnirp.org/documents/RFReview.pdf , accessed August 2009.

Kühn S, Lott U, Kramer A and Kuster N. Assessment of Human Exposure to

Electromagnetic Radiation from Wireless Devices in Home and Office Environments. Available at http://www.who.int/peh-emf/meetings/archive/bsw_kuster.pdf, Accessed August 2009.

Malone D, Duffy K, and Leith D. Modeling the 802.11 Distributed Coordination Function in Non-saturated Heterogeneous Conditions. *IEEE/ACM Transactions on Networking*, 15(1):159–172, February 2007.

Malone D and Malone LA. Ambient Radiofrequency Power: The Impact of the Number of Devices in a Wi-Fi Network. *Health Physics*, 96(6):629–635, June 2009.

Multiband Atheros driver for WiFi (MADWIFI) project. Available at http://sourceforge.net/projects/madwifi/. Accessed August 2009.

Navy Research Lab. Multi-Generator (MGEN) Version 4.2. Available at http://pf.itd.nrl.navy.mil/mgen/mgen.html. Accessed August 2009.

Azadeh P, Calderon C, Maan S, Khalid M, Addison D and Mee T. Evaluation of Exposure of School Children to Electromagnetic Fields from Wireless Computer Networks (Wi-Fi). BioEM2009.

Soekris Engineering. net4801 series Boards and Systems User's Manual. Available at http://www.soekris.com/manuals/net4801_manual.pdf. Accessed August 2009.

Schmid, G.; Preiner, P.; Lager, D.; Uberbacher, R.; and Georg, R. Exposure of the General Public due to Wireless LAN Applications in Public Places. Radiat Prot Dosimetry, 124(1):48–52; 2007.

Appendix

Here we show the details of the calculations for the radiated power in a number of situations. First, consider the upload of a file through a 1Mbps broadband link. We need to determine the total number of bytes per second and packets per second being sent over the network so that we can use equation (16). We note that there will actually be two senders in this wireless network: the station uploading the data and the access point, which will be sending higher level response packets. Protocols are usually designed so that these response packets are sent for every one or two packets sent, but will be much smaller (60 or 70 bytes). To estimate the number of packets per second that can be sent over a 1Mbps link, we need to know the packet size in bits. Packet sizes of 1400-1500 bytes are typical on modern broadband networks, so we use 1400 bytes = 11200 bits. This gives a figure of 1Mbps/11200 = 89 packets per second in one direction. We double this, to allow for the responses in the other direction. The number of bits per second will be 1Mbps in one direction and roughly 1Mbps * 70bytes/1400bytes = 0.05Mbps in the other direction. For a Wi-Fi rate of 11Mbps, we can use equation (16) to estimate the power as about 17mW. This is well below the saturated power of just over 80mW, so we do not need to make any adjustment.

As a second example, we consider our heavy YouTube user, who might spend long periods watching videos. This example is very similar to the previous example but the traffic now flows from AP to station. The constraint is how fast the user needs to download video in order to watch it for a period of time. Gill et al. 2007 show that most YouTube videos are encoded at a rate between 300 and 400Kbps, with the mean and median falling in this range. Starting with a rate of 400Kbps rather than 1Mbps, we may repeat the above calculation to get a value around 7mW.

Now, let us consider a classroom of 30 YouTube users. In this case we now have 400KB * 30 users traffic from the access point to the laptops, plus the response packets from the laptops to the access point. Calculating as above, we get a power estimate of around 203mW. Checking Fig. 9, we find that the summed duty cycle for ~30 saturated nodes is just over 1.2, suggesting that power actually saturates around 120mW.

As a final note, networks often contain a small amount of traffic that does not directly relate to higher-level user activity, which we have neglected in these calculations. For Wi-Fi networks, one source of this traffic is "beacon" packets, which advertise the network. These small packets are usually sent at a rate of 10 per second by the AP. Factoring in these packets results in a negligible change in the power.

Figures.

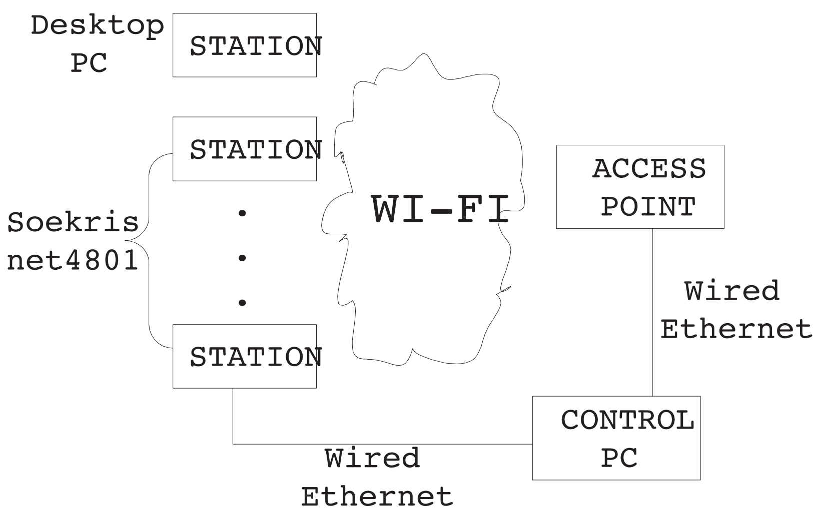

Figure 1: A Schematic of the Structure of Wireless Testbed.

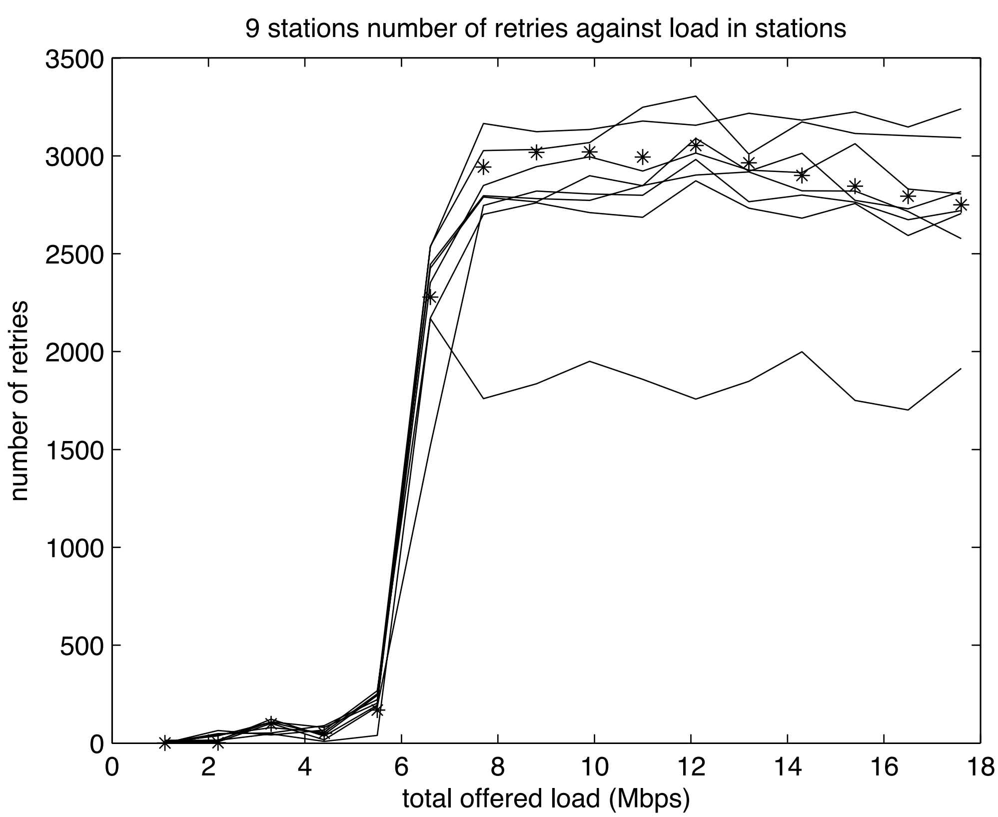

Figure 2: Number of Retries for each station vs. offered load. The dotted line shows the result for the desktop PC station, with the other 8 stations shown with solid lines. Big buffer experiment based on 11Mbps data rate and 100s experiment time.

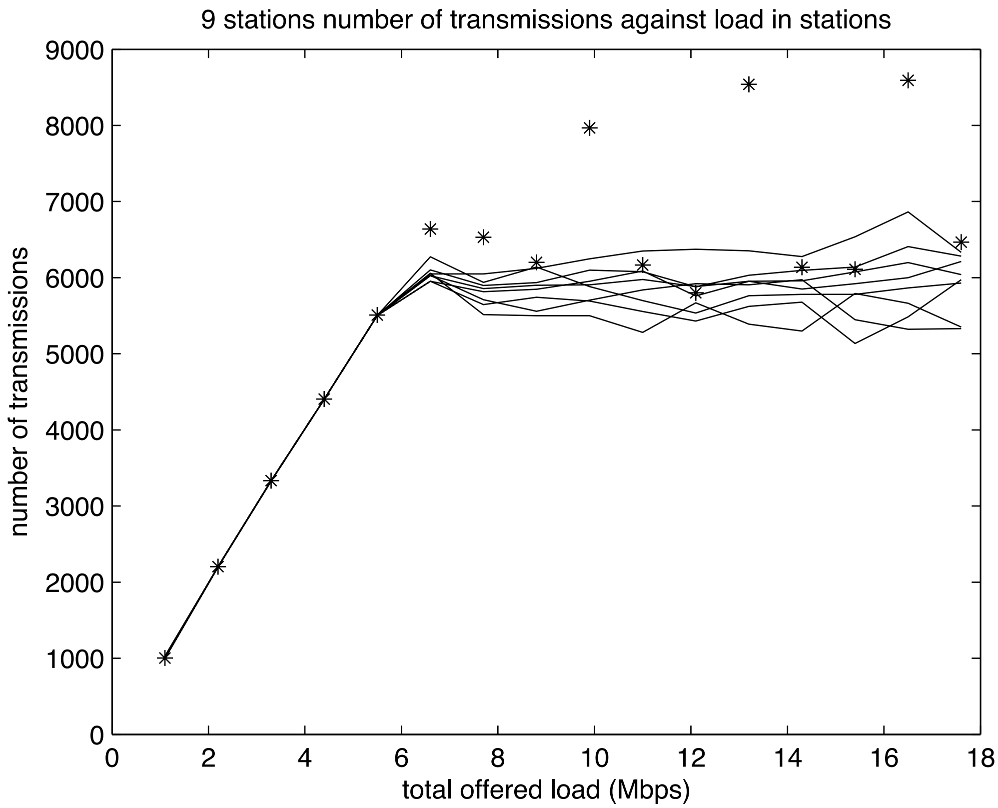

Figure 3: Number of Transmissions for each station vs. offered load. The dotted line shows the result for the desktop PC station, with the other 8 stations shown with solid lines. Big buffer experiment based on 11Mbps data rate and 100s experiment time.

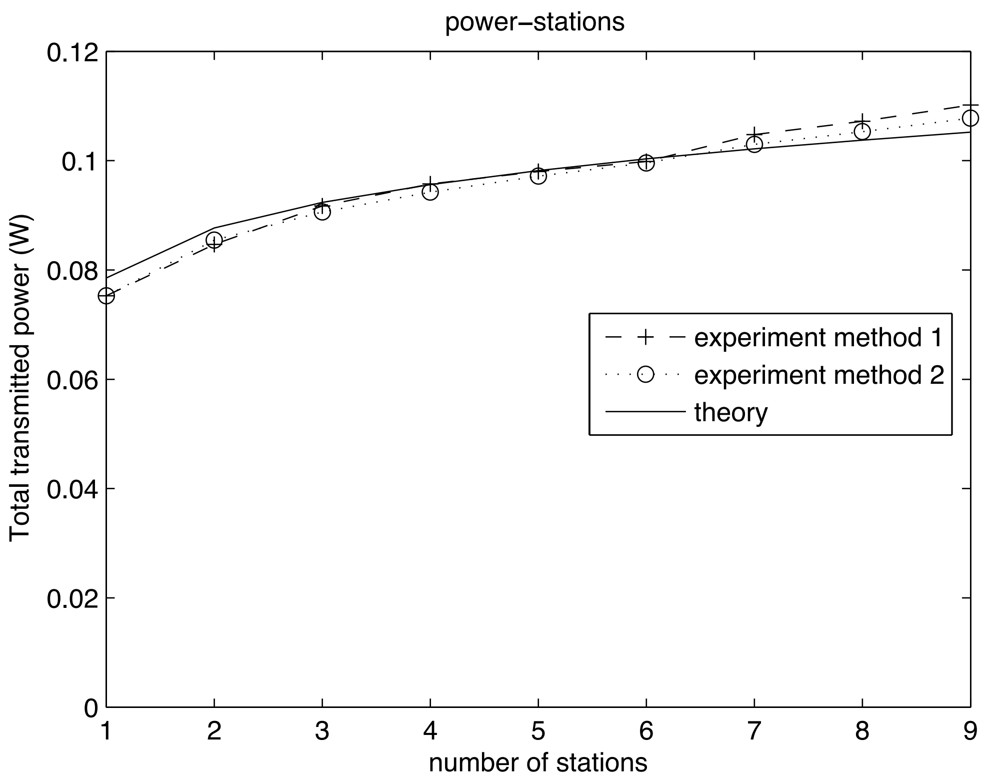

Figure 4: Transmitted power vs. Number of stations for saturated network. The nominal output of stations is 100mW. The result is based on 11Mbps data rate and 100s experiment time.

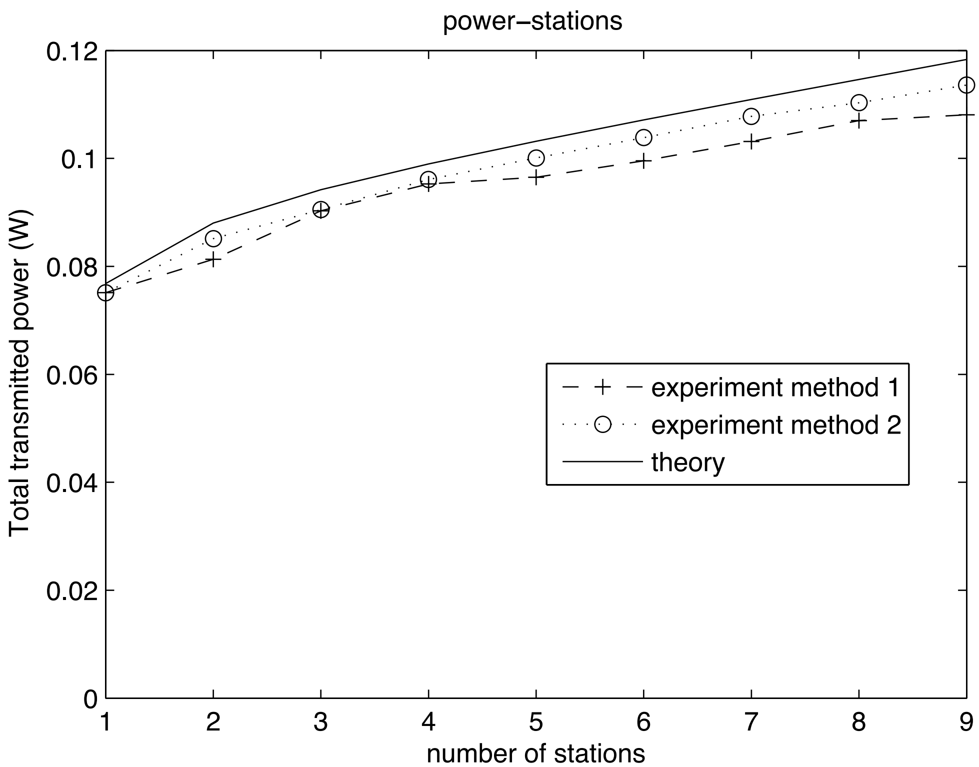

Figure 5: Transmitted power vs. Number of stations for saturated broadcast network. The nominal output of stations is 100mW. The result is based on 11Mbps data rate and 100s experiment time.

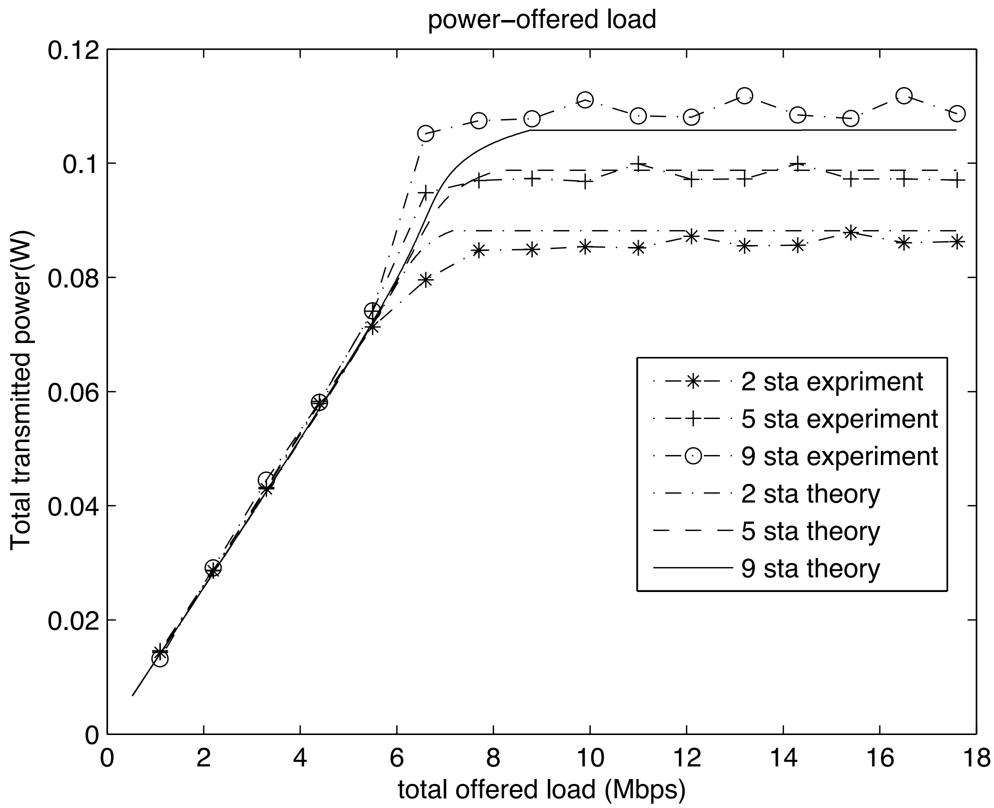

Figure 6: Transmitted power vs. offered load. Big buffer. The nominal output of stations is 100mW. The result is based on 11Mbps data rate and 100s experiment time.

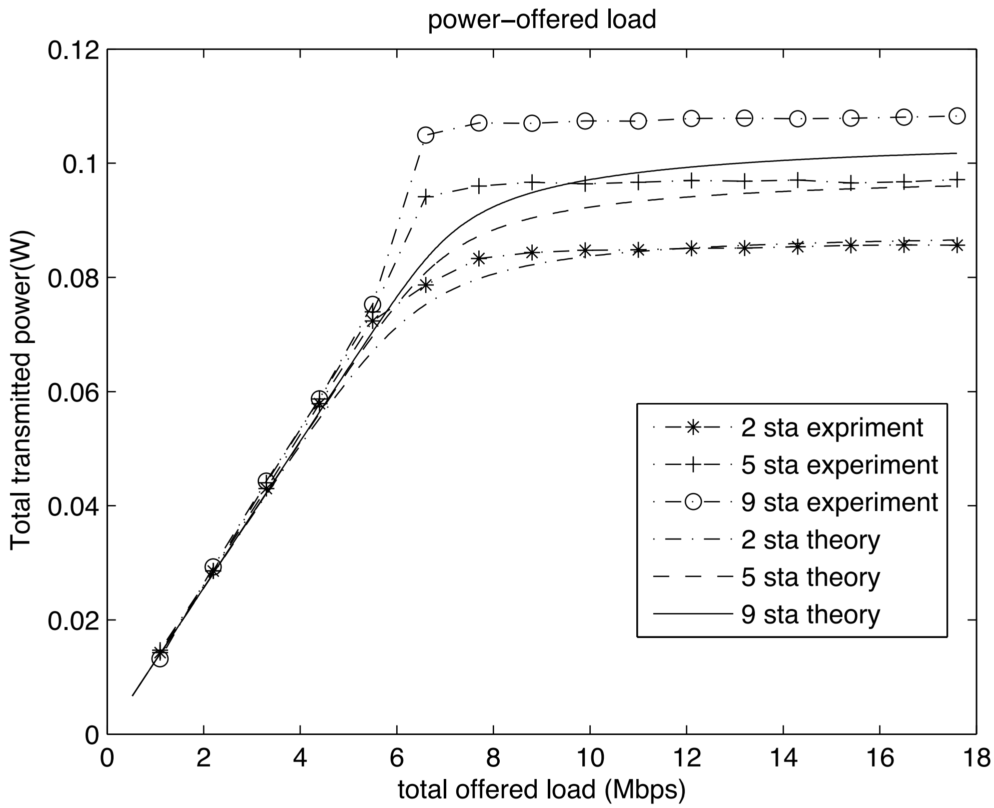

Figure 7: Transmitted power vs. offered load. No buffer. The nominal output of stations is 100mW. The result is based on 11Mbps data rate and 100s experiment time.

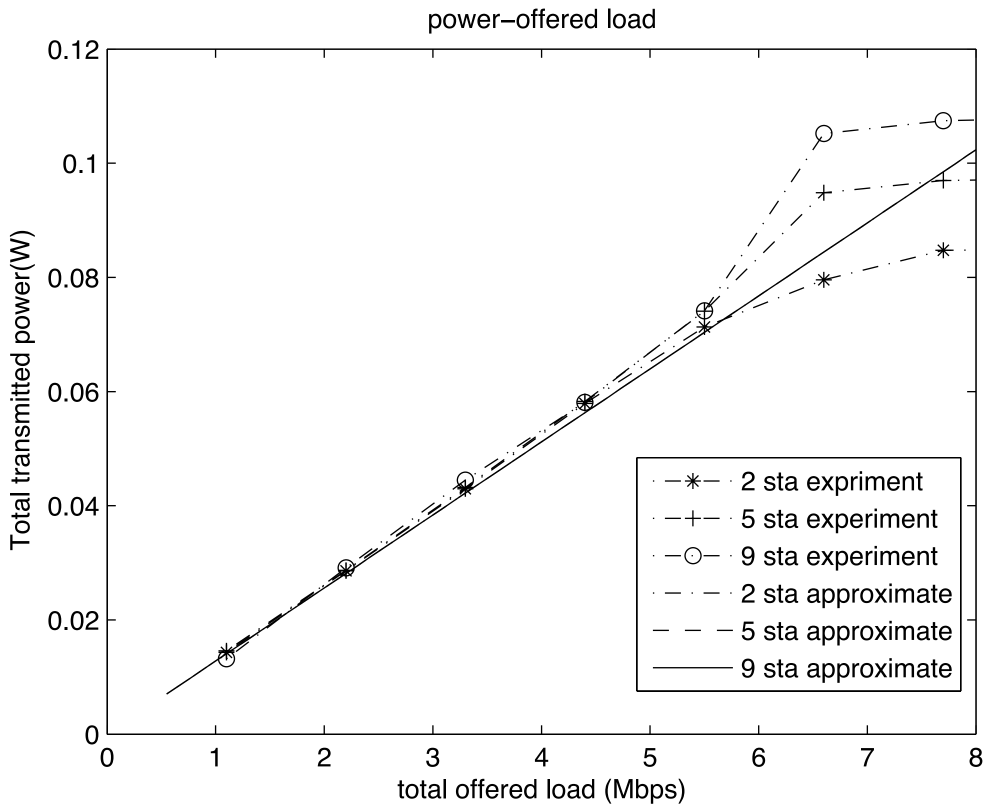

Figure 8: Power vs. offered load with simple lightly-loaded approximation. The nominal power of stations is 100mW. The result is based on 11Mbps data rate and 100s experiment time.

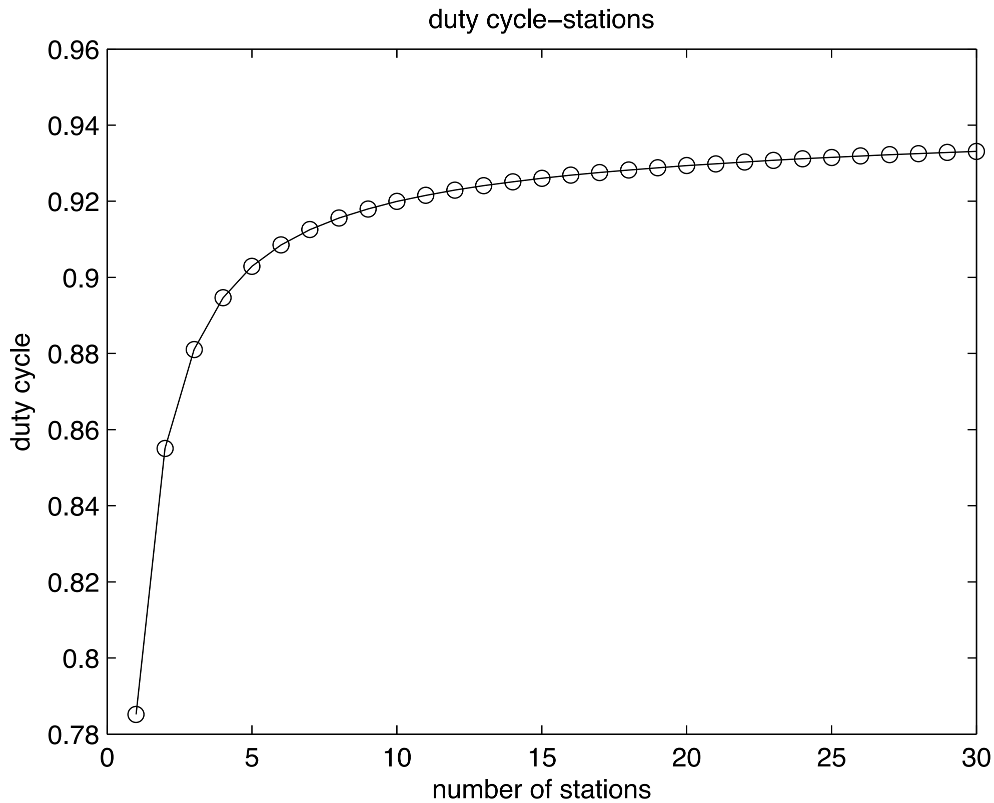

Figure 9: Network duty cycle and duty cycle summed over stations, as predicted by the model.

Figure 10: Duty Cycle vs. Number of stations for the saturated network. Based on 11Mbps data rate and 100s experiment time.

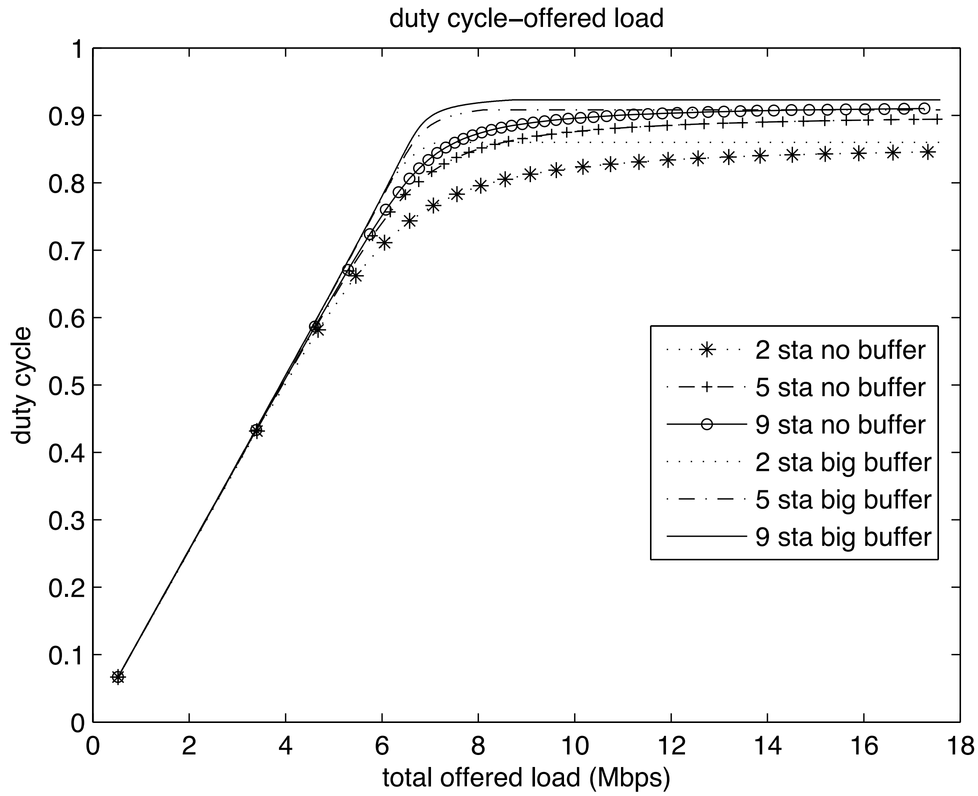

Figure 11: Duty Cycle vs. offered load for the non-saturated network. Based on 11Mbps data rate, big-buffer model and 100s experiment time.