Ambient Radiofrequency Power: the Impact of the Number of Devices in a Wi-Fi Network

David Malone[*]and Lesley A. Malone[†]

June 21, 2010

Abstract

This paper considers how the total radiofrequency power varies as the number of devices in a Wi-Fi network increases. Under the assumption that all devices in the network always have data to transmit, we show that work from network engineering can be adapted to calculate the total transmitted power, accounting for the possibility of multiple devices transmitting at the same time. The paper focuses on total transmitted power as it gives an upper bound on exposure and makes the impact of multiple devices clear. The results show that the number of devices does have a significant impact on the radiated power. For example, one station transmitting small packets in unicast mode gives about half the nominal power value, to reach nominal value around 25 devices are required, and 150% of the nominal value is achieved around 350 devices. We see that 802.11's transmission protocol is usually effective in limiting the power even in very large 802.11 networks.
This is a non-final version of the article published in final form in Health Physics. Vol 96, Issue 6, 629-635, June 2009. See `http://www.health-physics.com/`

exposure, radiofrequency; dose assessment; safety standards; computers

1 Introduction

Wireless Local Area Networks (WLANs) are now common in many places, including both homes and workplaces. These networks use variants of the IEEE (Institute of Electrical and Electronics Engineers) 802.11 standard, better known in the market as Wi-Fi. These devices transmit using relatively low power (typically 100mW or less) in the ISM (industrial, scientific and medical) band at 2.4 GHz or at 5GHz. There has been some interest in the likely doses arising from these networks, leading to a number of recent studies [10, 4].

One factor of note is that devices using the 802.11 standard do not transmit continously. Devices only transmit when they have data to send. Further, the 802.11 Medium Access Control (MAC) protocol regulates when transmissions may take place when data is available. This MAC protocol is designed so that (usually) transmissions take place one at a time, separated by small gaps.

Thus, the MAC protocol has an important impact on the power transmitted by an 802.11 network. For example, if we have n devices that all want to transmit, we might nïavely assume the power to be n times greater than if we have one device. On the other hand, if the MAC protocol was perfect, one could assume the power

to be exactly the same as if there was just one device. This second assumption is made in [10], where estimates are made assuming that only one device transmits at a time. The appendix of [4] gives an excellent discussion of factors that determine exposure from 802.11 networks, but only considers the case where one device can transmit at any given time.

The MAC protocol is randomised, and in practice sometimes multiple devices do transmit at the same time, resulting in a *collision*. Considerable effort has been made to understand collisions and the randomised MAC by the network engineering community. This is because there is a trade off between the collisions (which result in data being retransmitted) and the gaps between transmissions (during which no useful work is done).

In this paper, we will adapt some of the network engineering techniques to study how the total transmitted power scales up as the number of devices is increased in a Wi-Fi network. This is of particular interest to those who work in an environment with many wireless devices. A typical home may have one or two wireless devices whereas a class room, office or apartment block might well have 30 devices. A number of large-scale wireless testbeds have also been built, such as the ORBIT testbed [9], which is a grid of 400 wireless devices in a relatively small space. Indeed at some large technical conferences over 500 devices have been observed making transmissions [3].

Note that this paper focuses on the case where all devices in the Wi-Fi network always have packets to transmit. We also calculate the total transmitted power by the network (i.e. the sum of the power from all devices), rather than the exposure at a particular point. Thus we omit several important factors for calculating exposure, such as the distance from the devices, how much traffic they have to transmit and the pattern of absorption/reflection in the environment. These simplifications give upper bounds on exposure and allow us to focus on the impact of the number of devices.

2 Bianchi's 802.11 MAC Model

First, we will briefly describe the 802.11b MAC protocol [5]. When a device wants to transmit it picks a random number in the range 0 to CW, where CW is known as the contention window. Then it listens to see if there are other transmissions on the channel. While there are no transmissions it decreases its number every $20\mu s$. However, if a transmission is heard, this count down is suspended until the channel is idle for $50~\mu s$, when the count down resumes from where it left off. When a device's count down reaches zero it transmits.

If a transmission is successful, the device that receives the packet will send a response, called an ACK, letting the sender know that the data was successfully received. This ACK is sent $10~\mu s$ after the data is received, so no other devices will begin transmitting before the ACK is sent. After a success CW is set to its minimum value of $CW_{min} = 31$.

If the transmission is unsuccessful, usually because the data was corrupted by noise or another device transmits resulting in a collision, then no ACK is sent. In this case the sending station doubles its CW value (up to a maximum of $CW_{max} = 1023$ and tries again. After a certain number of bad tries, the data will be discarded and the device will move on to the next available data.

Bianchi [2] modelled a network of devices following this protocol using a mix of Markov chain [7] and more traditional network modelling techniques [1]. He observed that all devices moved through states in this backoff procedure at the same time and could estimate the probability that a station was about to enter a state where it would transmit, which we call τ. This probability is dependent on

the MAC parameters such CW_{min} and $CW_{max} = 1023$, and also the number of stations in the network n.

Bianchi makes certain simplifying assumptions. All stations are assumed to be identical and to have a limitless stream of packets of data to send. It is also assumed that transmission failures only occur because of transmissions and that stations will retransmit each packet until it is successfully sent. These assumptions have been relaxed by subsequent research, however for simplicity we will begin with these assumptions.

3   Calculating Transmitted Power

In the Bianchi model, the amount of data successfully transmitted (the system throughput) is derived by calculating the average amount of data transmitted by the network as it moves through these Markov states and the average amount of time taken to move through the states. We will now calculate the average power transmitted by the system in an analogous way.

First, let us calculate energy transmitted per state. If no stations transmit, then no energy is transmitted. If a single station transmits, then we will have a single transmission followed by an ACK. If $r \geq 2$ stations transmit, then we will have r simultaneous transmissions followed by a timeout while the transmissions wait for an ACK, but do not receive one. Thus we write the mean energy per state as the sum of these three contribution:

$$E = 0(1-\tau)^n + E_S n\tau(1-\tau)^{n-1} + E_C \sum_{r=2}^{n} r\binom{n}{r}\tau^r(1-\tau)^{n-r}, \qquad (1)$$

where E_S is the mean energy radiated in successful transmission (including the ACK) and E_C is the mean energy radiated per-station in a collision. The expression for E may be simplified as

$$(E_S - E_c)n\tau(1-\tau)^{n-1} + E_C n\tau. \qquad (2)$$

The mean length of a state is calculated in the same way as in the Bianchi model:

$$T = T_I(1-\tau)^n + T_S n\tau(1-\tau)^{n-1} + T_C(1-(1-\tau)^n - n\tau(1-\tau)^{n-1}), \quad (3)$$

where T_I is the length of an idle state, T_S is the mean time for a successful transmission and T_C is the mean time for the length of a transmission. The mean power is then:

$$P = \frac{E}{T} = \frac{(E_S - E_c)n\tau(1-\tau)^{n-1} + E_C n\tau}{T_I(1-\tau)^n + T_S n\tau(1-\tau)^{n-1} + T_C(1-(1-\tau)^n - n\tau(1-\tau)^{n-1})}. \quad (4)$$

The times T_I, T_S and T_C can easily be calculated from the 802.11 standards by adding up the length of time required to transmit headers, data, ACK packets and also the various inter-packet times. We may also produce estimates of E_S and E_C by using the nominal power output (say 100mW, the maximum equivalent isotropic radiated power permitted in Europe) and the length of time actually spent transmitting headers, data or ACKs.

4   Results

Using the model from Section 3 we may predict how the power transmitted will scale in a busy 802.11 network as the number of stations is increased. Also, by
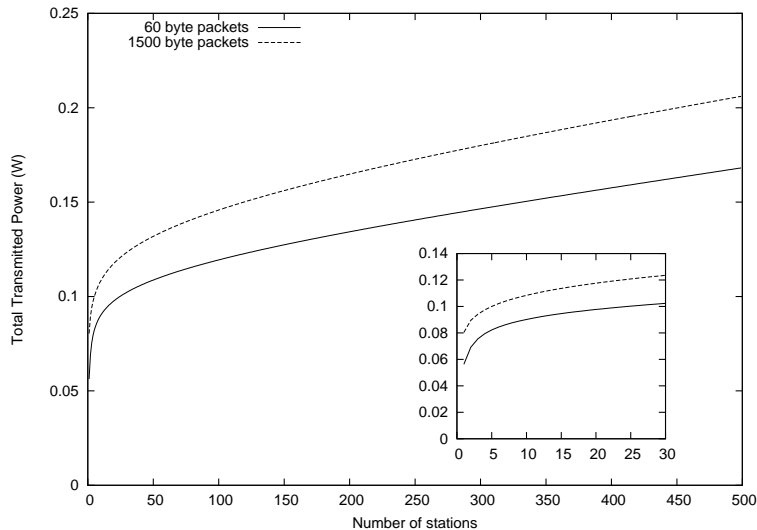
Figure 1: Power vs. number of stations. The nominal output of the stations is 100mW. The inset graph shows an enlarged version for small numbers of stations.

| | slope for 1–10 stations | | slope for 100–500 stations | |
	small packets	large packets	small packets	large packets
Unicast	3mW station^{-1}	3mW station^{-1}	0.1mW station^{-1}	0.1mW station^{-1}
Broadcast	6mW station^{-1}	5mW station^{-1}	5mW station^{-1}	6mW station^{-1}

Table 1: The approximate additional power per station in different situations.

considering variants of the model we may consider the impact of network errors or different traffic types.

Fig. 1 shows the results of this calculation for typical 802.11b settings with a nominal power of 100mW. We show the results for small packets (60 bytes) and large packets (1500 bytes). As packets become larger we naturally expect the power to increase as a greater fraction of the time is spent transmitting rather than on count down.

We note that the MAC actually keeps the transmitted power quite low. For the smaller size of packet and one station, the power is about half the nominal value. To get up to 100mW we need 25 stations in the network, and to get to 150mW we need around 350 stations. For larger packets the absolute values are higher: with one station we transmit about 80mW. However, it still takes 120 stations to get to 150mW and 450 stations to reach 200mW.

The *unicast* line of Table 1 shows the results of fitting two lines to each curve in Fig. 1 and taking the slope. The first line is fitted between stations one to ten, and we see that each station brings an additional power of about 3mW, regardless of packet size. The second line is fitted between stations 100–500. Over this range each station brings only 0.1mW.

In reference [6] Bianchi's model was extended to account for packets being dropped after a fixed number of bad tries. This work also allows the possibility that transmissions may be unsuccessful because of noise on the channel that is not due to collisions. Fig. 2 shows the results of adapting eqn (4) to be used with the model from [6].

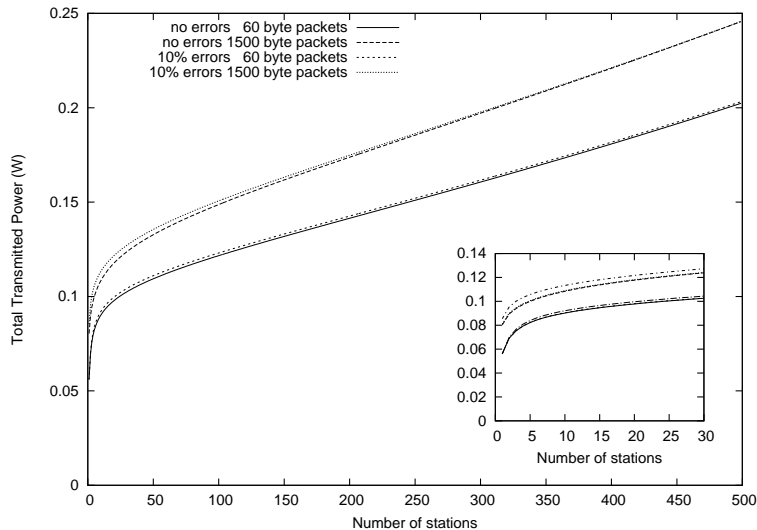We see that the power transmitted still increases slowly, though more quickly

4

Figure 2: Power vs. number of stations with a maximum of seven retries per packet and some packets lost due to noise. The nominal output of the stations is 100mW. The inset graph shows an enlarged version for small numbers of stations.

than in the case where packet retransmissions are repeated until successful. This is because discarding a packet resets the CW value to CW_{min}, and so transmissions are more frequent. However, the MAC still does a good job at limiting concurrent transmissions, and we require around 350 stations before the transmitted power reaches twice the nominal value. We also observe that the errors introduced due to noise have little impact on the transmitted power. In fact, a more realistic error rate of 1produces a change in power that is, for practical purposes, unchanged from the situation with no errors.

So far we have looked at models of *unicast* packets, i.e packets directed at some specific device in the network. This is the usual type of traffic found in a network. However, there is another type of traffic that can be directed to groups of devices in the network, called *broadcast* or *multicast*. Since there is no one device that can send an ACK for this traffic, it is unacknowledged and assumed to be successful. This means that the usual backoff mechanism does not function for broadcast packets.

Usually there is only a small amount of broadcast traffic in a WLAN, and so it should not have a big impact. Fig. 3 shows the extreme case where all stations send a constant stream of broadcasts. We see that the lack of the backoff mechanism results in significantly more power being transmitted: 500 stations can now transmit almost 3W. The *broadcast* line of Table 1 shows the fitted slope for stations 1–10 and stations 100–500, we see that in this situation each station brings about additional 6mW. When compared to the unicast traffic above, we see that the backoff meachnism for unicast traffic is actually quite effective in controlling the increase in transmitted power.

Finally, we look at the impact of using the newer 802.11g standard, which offers speeds of up to 54Mbps, compared to 802.11bs 11Mbps. There are a number of different modes in which 802.11g can operate: it has a backward compatible mode that interoperates with 802.11b and a faster mode that can be used when all devices are 802.11g capable. In addition to the faster data rate, a number of the overheads are smaller in 802.11g (for example, counter decrements can happen every 9s and packet preambles are shorter). Fig. 4 shows results for a pure 802.11g network. For
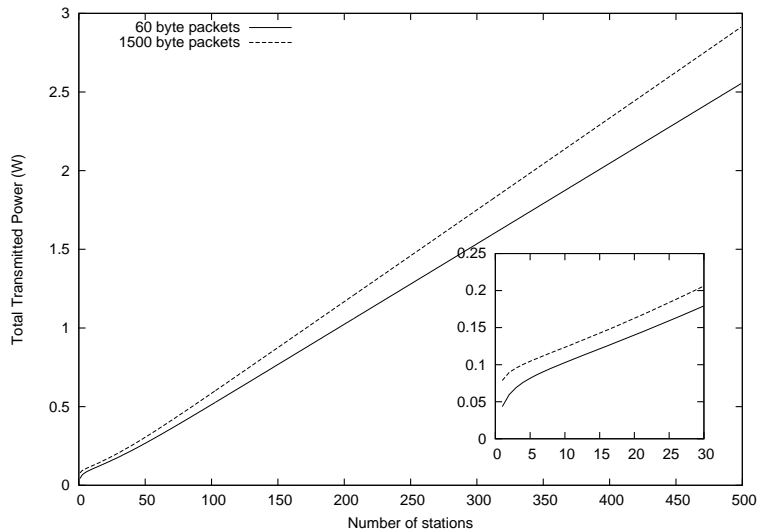
Figure 3: Power vs. number of stations when all transmissions are broadcasts. The nominal output of the stations is 100mW. The inset graph shows an enlarged version for small numbers of stations.

large packets, overall trends are very similar to the 802.11b results in Fig. 1, though power output is slightly decreased because of the shorter packets. For small packets, the reduction in power output is roughly half that of 11b, because the reduction in size of the preamble is a more substantial fraction of the entire packet.

5 Discussion

As noted in the introduction, we have ignored many factors here and concentrated on the total power transmitted by an 802.11 network. One factor of note is that we have assumed that it is a network operating on a single channel. In fact, there are multiple channels available for 802.11 use, which can operate independently in parallel. This means that we can calculate the power on each channel and sum them. For 802.11b/g there are effectively three independent channels available. The 802.11a allows the use of more channels in the 5GHz band, but these bands are less frequently used.

The International Commission on Non-Ionizing Radiation Protection (ICNIRP) guidelines [8] for exposure in this frequency range is 80mW kg^{-1} for the general public and 400mW kg^{-1} for occupational exposure (whole body). To reach these levels of exposure a person would need to absorb a substantial fraction of the energy emitted from a large number of stations. It seems unlikely that this could be achieved in normal circumstances: for small networks the power levels are too low when the exposure is spread over more than a couple of kilogrammes of tissue; in large networks the physical spacing of the devices should reduce the power density considerably.

It is worth noting how the shape of the power curves depends on the values of certain parameters. For example, if stations wait a very long time for an ACK after a transmission, then it is possible for the power transmitted to initially decrease as stations are added, because collisions with small numbers of stations will then have a lower power than a successful transmission. However, the long term trends are the same: there is an initial balancing of good transmissions against collisions
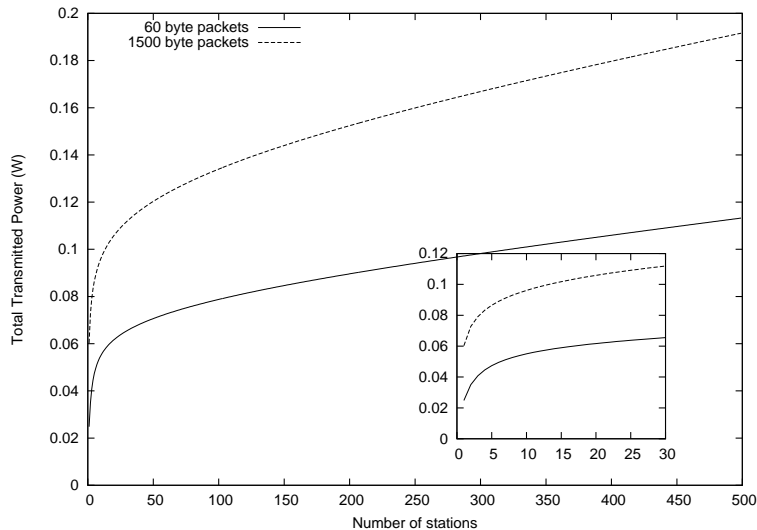
Figure 4: Power vs. number of stations for pure 802.11g network. The nominal output of the stations is 100mW. The inset graph shows an enlarged version for small numbers of stations.

followed by a gradual increase in power as collisions involving larger numbers of stations begin to dominate the network.

Where this long-term trend sets in will depend on the CW_{min} and CW_{max} parameters. These values are fixed for 802.11b, but can be changed in an ammendment called 802.11e. This allows certain data packets to be transmitted with larger or smaller values, in order to prioritise or deprioritise them. While supported by many pieces of hardware, this extension is not yet in wide practical use.

Overall, we have seen that the 802.11 MAC is quite effective in keeping the number of simultaneous transmissions low. Indeed, adding stations to an already busy network results in a relatively small increase in the power in a unicast network. We can see that in moving from from one station to three we see a similar increase in power as moving from 100 to 300. This may be useful to those aiming to control total power output, as this indicates that only substantial changes in numbers of devices are likely to have a significant impact on overall power.

References

[1] D. P. Bertsekas and R. G. Gallagher. *Data Networks*. Longman Higher Education, 1986.

[2] G. Bianchi. Performance analysis of IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, March 2000.

[3] Andrea G. Forte, Sangho Shin, and Henning Schulzrinne. IEEE 802.11 in the large: Observations at an IETF meeting. Technical report, Columbia University, 2006.

[4] Kenneth R. Foster. Radiofrequency exposure from wireless LANs utilizing Wi-Fi technology. *Health Physics*, 92(3):280–289, March 2007.

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

[6] Q. Ni, T. Li, T. Turletti, and Y. Xiao. Saturation throughput analysis of error-prone 802.11 wireless networks. *Wireless Communications and Mobile Computing*, 5(8):945–956, 2005.

[7] J. R. Norris. *Markov Chains*. Cambridge University Press, 1998.

[8] 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.

[9] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh. Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols. In *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, 2005.

[10] G. Schmid, P. Preiner, D. Lager, R. Überbacher, and R. Georg. Exposure of the general public due to wireless LAN applications in public places. *Radiation Protection Dosimetry*, 124(1):48–52, 2007.