Performance Monitoring of a wireless LAN and some design heuristics

By Sanjeev Dwivedi and Anandha Gopalan

Abstract

As the popularity of the wireless LANs increases, the design and analysis of a wireless LAN become more and more important. These problems are not very well studied and hence often the inappropriate design choices lead to a shaky network. In this report we summarize the results of experimenting with a wireless network and observing the effect of congestion on various parameters affecting the network performance.

Introduction

Wireless Local Area Networks are a relatively new concept and how to design and install a Wireless Network is a difficult task. It becomes especially difficult in wake of the fact that the problems associated with the wireless environment are not very well studied. Also, the network behavior when the load increases has not been studied very well.

Various ways to analyze and design the performance evaluation in a wireless network are:

(1) Mathematical Analysis
(2) Computer Simulation, and
(3) Performance monitoring on a test bed.

The data obtained from mathematical analysis and computer simulation are usually purely speculative since the characteristics of signal propagation, are not very well known. Hence it is most beneficial to perform the experiments on a real Wireless LAN to obtain the actual data.

In this report we summarize the results of the experiments we performed in a wireless local area network. The goal was to measure the parameters of the network, which will give us some insight towards the actual design of a system and help us understand the effect of increasing load on the throughput of the network. This, we expect will give us an insight into the behavior of the network and hence will allow us to design more effective wireless networks. It will allow us to gather more information on how much infrastructure is actually required to maintain an effective wireless network that will give us the same benefits (or at least comparable) to the benefits offered by the wired network and hence making the transit from wired to wireless network less painful.

We expect that the network performance should go down appreciably as more and more users come in because of the fact that the different users will be contending for the same fixed amount of resources and also because the structure of the CSMA/CA protocol (based on the IEEE 802.11) that we have used for our experiments the, delay period i.e. the delay before which a mobile waiting to transmit data will get the channel to transmit data will increase hence affecting the network adversely. What is not known is how dramatically will this actually affect the network.

Experiment methodology and the test bed

The infrastructure that we used for the experiments were:

- Wavelan's AP-1000 (11 Mbps) Access point
- Eight Clio Handheld devices running Microsoft Windows CE operating System equipped with wavelan silver cards (using IEEE 802.11)
- One HP-Jornada running Microsoft Windows CE Operating System using the same wavelan silver card.
- One Laptop Running RedHat Linux 6.1
- One Laptop Running Windows 2000

Wavelan Silver Cards (IEEE 802.11 compliant) were used in all these mobile devices to perform the experiments.

We have used the Orinoco Client Manager supplied with the Wavelan cards for the signal to noise measurements. For experiment 1 described below, packet transmission and reception detection has been preformed with the Link test utility present with the Orinoco Client manager.

For experiment 2 we have used the TFTP (trivial ftp) application to transfer the files

For battery power measurements we have used the HP-Battery manager supplied with the HP-Laptop.

Experiments

Experiment 1:

Measuring the packet error rate with varying signal to noise ratio and varying load on the network.

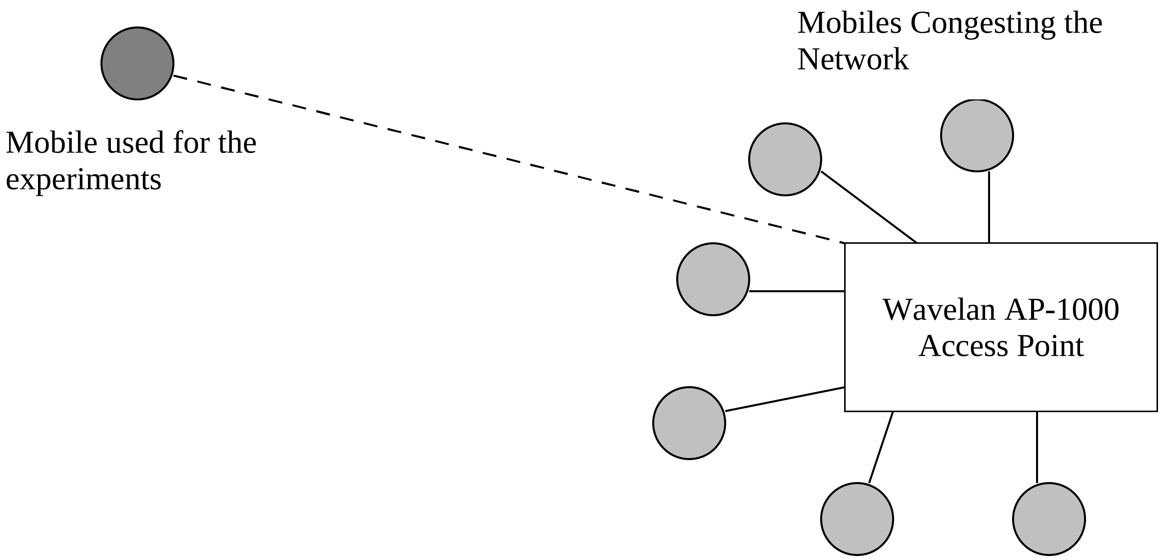

We used one of the mobile devices to perform the measurements while the other devices were congesting the network. To increase the signal to noise ratio, we moved the mobile device away from the access point. In the various runs of the experiment we kept changing the number of devices congesting the network and then using the other mobile to perform the measurements with varying signal to noise ratios.

Using the same setup as above we measured the battery power consumed in transmitting a fixed number of packets as the congestion increased. We found out that the battery power consumed for only transmission does not change irrespective of the congestion in the network.

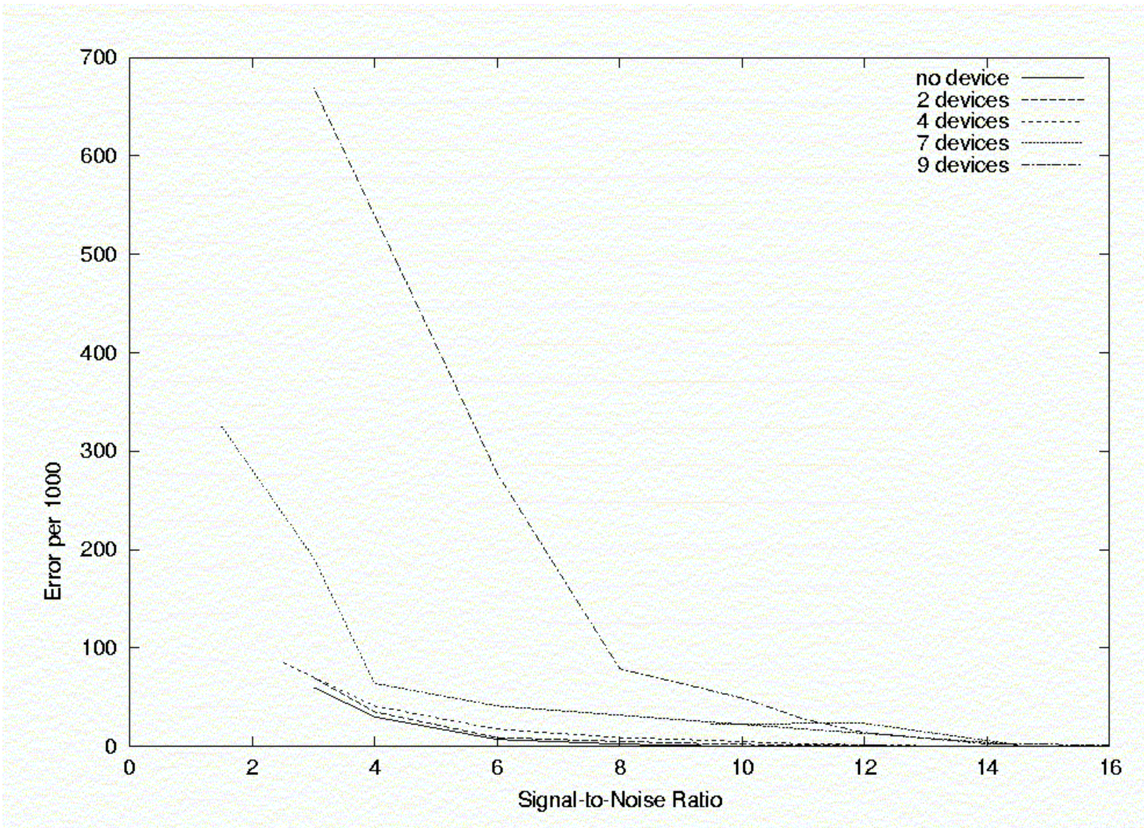

Experiment 2:

Measuring the throughput of the network as the congestion goes up.

We had the same experimental setup and we tried to measure the time taken for a file to be successfully transmitted as the congestion in the network went up. We also measured the number of retries it took for the file to go through. Since any TCP based protocol timeouts after some retries and the semantics of TCP start affecting the network we used a UDP based file transfer protocol to transfer the file.

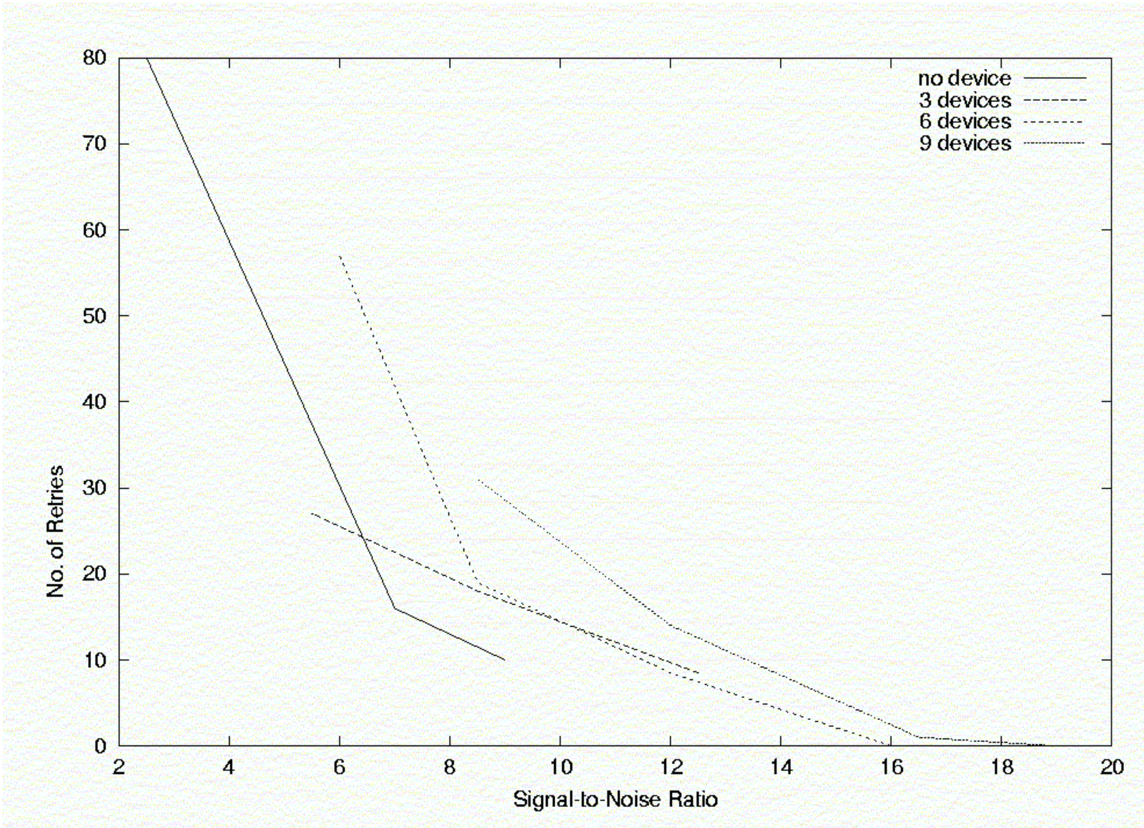

no device
3 devices
6 devices
9 devices
No. of Retries
Signal-to-Noise Ratio

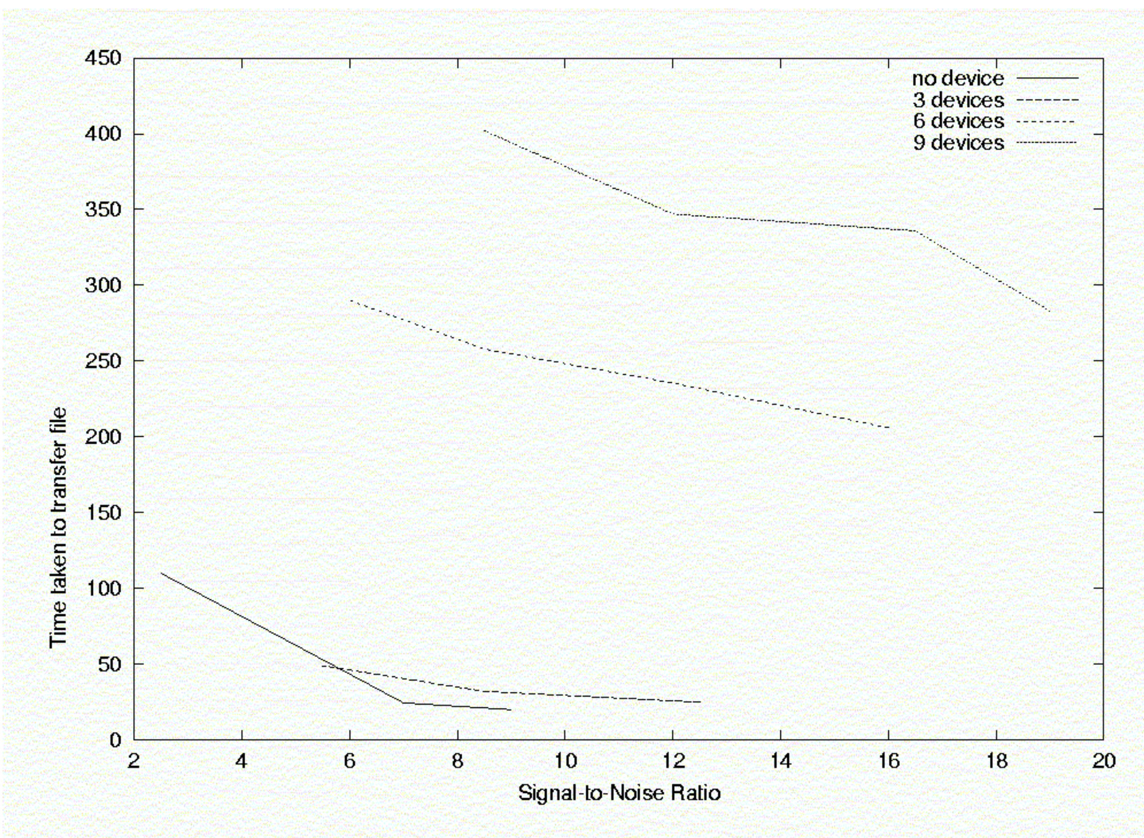

no device
3 devices
6 devices
9 devices
Time taken to transfer file
Signal-to-Noise Ratio

Conclusions

Experiment 1:

From experiment 1 we conclude that as the congestion in the network increases, the packet error rate increases. But the main point to notice is that as the network reaches a threshold the packet error rate increases very sharply and the performance of the network goes down appreciably. This point is important because while designing the network because in scenarios as shown below, this result may provide important applications:

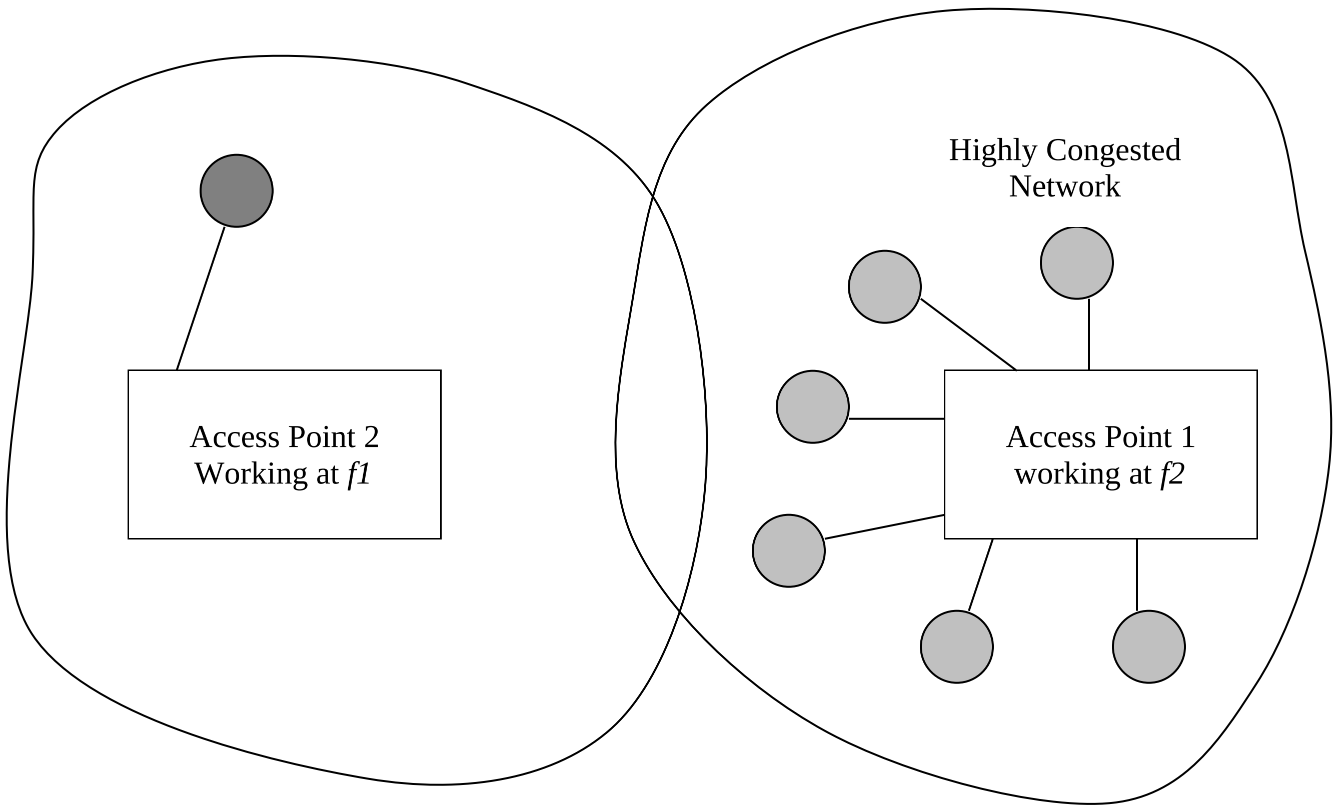

An example Wireless Network

In the above framework we see that Access Point 1 (operating at frequency *f1*), is serving a very large number of clients and the network around it is highly congested (probably working in the region where the network is almost at the point of collapse) whereas Access Point 2 (operating at frequency *f2*) is having very little traffic since there is only one mobile client using it. If at this point of time we shift the frequency f2 also to Access Point 1 the network will come in the region where the congestion will be less than threshold and hence the behavior of the network will be pretty good. The client who previously was in the region for the Access Point 2 might face some degradation in performance but that will not be substantial because it will have to use Access Point 1 where the network is not congested. We know that in an uncongested network the performance remains good even when the signal to noise ratio goes down quite a lot. So the single mobile, which was previously in the network of Access Point 2, will be performing reasonably while the clients of Access Point 1 will get improved performance.

Experiment 2:

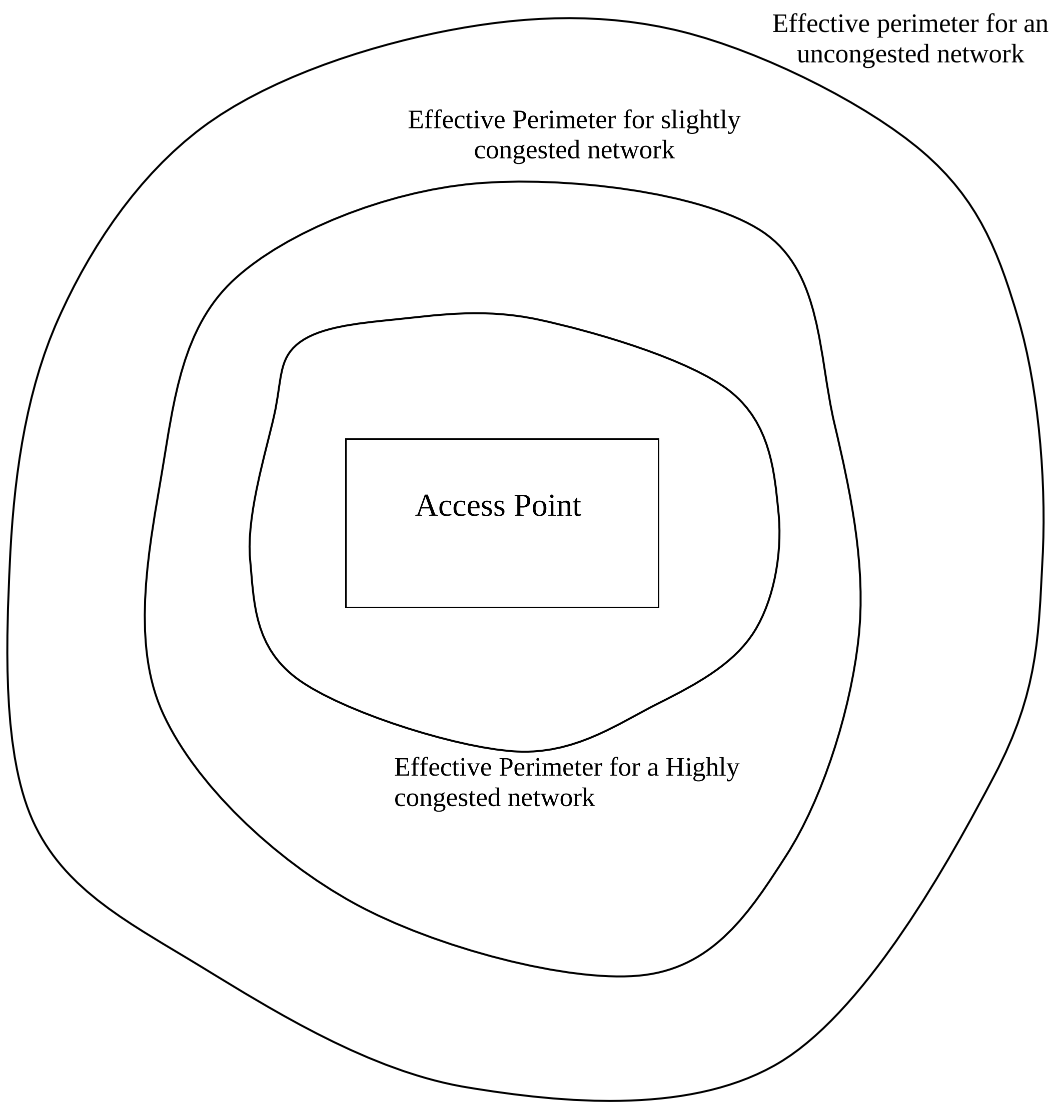

Effect of congestion on the range of an access point

From experiment 2 we conclude that as the congestion in the network goes up the effective perimeter of the network goes down and hence for wireless local area networks, the number of access points should be pretty large and located close to each other. This has another implication in the mobile. If there were some way in which the access point can send more information though its beacon regarding the traffic on the network, then based on the statistics, that we have compiled, the mobile can decide which access point it can use, thus making sure that they user has a good throughput.

Future work

(1) Having a more reliable means of measuring the throughput
(2) We do not have sufficient confidence in the experiments because we believe that the experimental setup we used to congest the network does not actually congest the network a lot. The transmission and reception characteristics of these devices are severely limited by their architecture. For example the CPU of these devices are very slow (typically in the range of 16-30Mhz), the bus speed, the memory etc are also very limited. Hence to get more confidence in our readings we would have liked to have devices, which don't have these limitations (e.g. laptops with more efficient processing units, memory etc.)
(3) Actually measuring the diversion of the frequency between the access points and finding out the effect.

References

IEEE 802.11 Wireless Local Area Networks, Crow et. all, IEEE Communications Magazine, September 1997.

Wavelan, http://www.wavelan.com