Reading Aguayo et al.

Philip Levis
April 3, 2011

The reading for Friday's class is a seminal wireless research paper. What makes this paper especially interesting experimentally is that while it is astoundingly careful, deep, and complete, ultimately one of the major conclusions of the paper turns out to be incorrect. Furthermore, this error is due to a mistake in the experimental methodology. This document explains the error and provides enough of a background in wireless to (hopefully) understand it. Many of the observations in the paper are correct: it's just one, albeit a major one, that is not.

It's worth stating up front that this mistake shouldn't be seen as a failure of the authors -- no one has before or since done a better job; they were (and still are) top researchers in the field renowned for their precision and care. But at the time the community was figuring out how WiFi networks behave in the field.This paper definitively established that many of the simplifying assumptions of wireless researchers had used can lead to grossly incorrect conclusions.

Signal to Noise Ratio

While not completely accurate, sound is a good approximate analogy for wireless communication. The inaccuracies are due to how we tend to encode data in sound as well as how fast sound travels compared to electromagnetic waves (hundreds of feet a second versus approximately a billion feet a second), but at a high level they behave similarly.

One of the general goals in wireless is to maximize the throughput -- bits of data delivered per second -- between a pair of devices. Where the paper talks about "link layer" measurements, this is exactly what they mean: measurements of how the link, or single hop communication, between two wireless nodes behaves.

There is a theoretical upper bound on the throughput of a link, called the Shannon limit (from Claude Shannon's information theory). The bound is based on the signal-to-noise ratio (SNR). Returning to the sound analogy, the signal strength is the volume of the signal you're trying to understand, while the noise is the signal strength of everything else. For example, a conversation in a quiet room (very low noise) can have a high signal to noise ratio (easy to understand), even if the other person is speaking softly (has a weak signal). In contrast, it can be hard to hear someone (low signal to noise ratio) when they're shouting (high signal) at a concert, because of the loud music (high noise).

Some more advanced analyses call it the signal to interference and noise ratio (SINR), separating the noise term into hardware noise, which relates to energy from within the receiver itself, and interference, energy from other signals in the environment.

Signal processing systems typically measure signal and noise in terms of energy, decibels (dB) or millidecibels (dBm). Decibels are a logarithmic scale: 10dB is ten times stronger than 0dB, and

20dB is 100 times stronger than 0dB. 0 dB corresponds to one watt, while 0 dBm corresponds to one milliwatt. WiFi transmits at about 20-23 dBm (100-300 mW), and typical hardware noise is in the range of -100 to -95 dBm. Because decibels are logarithmic, the signal to noise ratio is simply the strength of the signal minus the strength of the noise: a signal of -65 dBm and a noise of -95dBm has a signal to noise ratio of 30dB.

A higher SNR means that you can receive data at a higher rate (more bits per second). A lower SNR correspondingly lowers the maximum data rate. These are theoretical limits: practical systems typically do not achieve them. For example 802.11b, which the paper studies, has four possible rates: 1, 2, 5.5, and 11 Mbps. Operating at 5.5 Mbps requires a higher signal to noise ratio than 11 Mbps, but increasing the signal to noise ratio past what 11 Mbps requires does not give any benefit, as the actual wireless technology cannot take advantage of it. Figure 12 of the paper, shown at right. measures how the percentage of received packets changes as SNR increases.

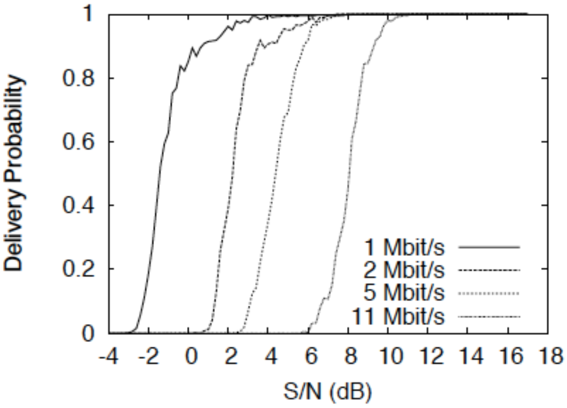

One implication of having different bitrates is that going to a lower bitrate can actually improve throughput. Returning to the sound analogy, bitrates correspond to how quickly you're speaking, the rate at which you are trying to transmit data. Sometimes, if it's noisy and you speak too quickly (too high a bitrate), the other person can't understand anything. But if you slow down and speak more slowly, they can. Choosing a lower bitrate increases the actual data transmission rate. Picking a bitrate that's too high for a given signal to noise ratio reduces throughput, because the receiver can't understand what's being sent. But picking a bitrate that's too low also reduces throughput, because you could speak more quickly and still be understood.

Measurement Flaw

All of the above description makes an important simplifying assumption: the signal to noise ratio is constant (or at least varies very little) for the duration of a packet. This is because of the way packets are encoded: if one part of the packet is decoded improperly (errors in recovering the bits), the entire packet fails.

If there aren't interfering transmissions, the signal to noise ratio is reasonably stable. The noise value is very stable: it's basically affected only by temperature, so while it might change over long time periods, it typically does not change much over the range of minutes.

Signal strength can change pretty quickly: put your hand between your wireless card and the AP. and you'll see WiFi signal strength drop by 20dB (99%) or more. These changes, while fast to our perception, though, are not very fast in terms of how long it takes to transmit a packet. Take a back of the envelope calculation: a 1.5kB packet, the standard WiFi packet size, is 12,000 bits. At 11Mbps, this takes just over a millisecond to transmit. It turns out in practice to take a little bit longer for some low-level reasons (the beginning of a packet is actually always at 1Mbps), but even then, 1.3ms is a conservative estimate. The sorts of environmental changes that affect signal strength don't tend to be this fast. If the Roofnet nodes were in fast-moving cars, it might be a bit different, but they're not: they are stationary.

Interference, in contrast, can change very quickly: interfering packets can be shorter than a millisecond, can appear spontaneously, and there might be very little regular pattern to their appearance.

This matters tremendously because of how the paper measures the signal to noise ratio. Section 2.2 reads:

> The Prism 2.5 chip-set provides per-frame measurements called RSSI (receive signal strength indication) and "silence value." The RSSI reflects the total power observed by the radio hardware while receiving the frame, including signal, interference, and background noise. The silence value reflects the total power observed just before the start of the frame. We found that the accuracy of the RSSI and silence readings was within 4 dB by comparison with a spectrum analyzer. This paper reports signal-to-noise ratios derived from the RSSI and silence values.

What this means is the "noise" value they use when computing the signal to noise ratio is the RF energy the radio hears immediately before the packet (e.g., the sound before someone talks, not while they talk). This energy could be due to the hardware noise of the radio, or interference from other nodes. This introduces two problems: measurement bias and imprecision.

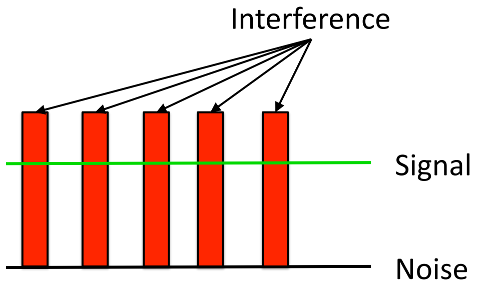

Imagine a link, shown on the left, that has a low hardware noise (-100dBm) and good signal (-70dBm). All packets that are successfully received will reflect a high signal to noise ratio of 30dB. Now suppose there is a nearby strong interfering transmitter that prevents reception while it is active. Suppose this interferer transmits 50% of the time. For the packets it receives, the receiver will see a signal to noise ratio of 30dB, suggesting perfect reception, but it will see a packet reception ratio of 50%.

The issue is that the noise measurements are only for received packets: if the "noise" value is changing quickly, this introduces a measurement bias, where the measurements are only for packets with the low noise value.

Furthermore, this measurement is imprecise. What matters is the noise during a packet, yet the silence measurement is taken before a packet. In practice, this means that you can have nonsensical signal to noise ratios, such as -20dB. Very low values like this occur when the interference ends just before the packet: the silence value can be very high, yet during the packet transmission there is no interference and the node receives the packet successfully.

Why This Matters

The reason this measurement flaw matters is it leads the paper towards the conclusion that multipath is a major contributor to strange link behavior. In wireless, multipath refers to the problem of the same signal arriving through different paths with slightly different delays. Going back to the sound analogy, it's like hearing someone shout in a canyon, such that all of the echoes (multiple paths) make it hard to understand what's said. Multipath does exist in networks like Roofnet, but it unlikely to be the principal cause of the observed link dynamics. We'll go over in class some evidence in the paper that suggest so.

Question

The paper examines the hypothesis that the low packet reception ratios observed might be due to interference (Section 8), and discards it. Why might their experimental setup and analysis lead to the wrong conclusion? What might you have done differently?