

Building a Wireless Network for a High Density of Users

David Lang - Intuit

Abstract:

Why do conference and school wireless networks always work so poorly? As IT professionals we are used to the network 'just working' and fixing things by changing configuration files. This mind-set, combined with obvious-but-wrong choices in laying out a wireless network frequently lead to a network that seems to work when it's tested, but that then becomes unusable when placed under load. This is at its worst at technical conferences where there are so many people, each carrying several devices, all trying to use the network at the same time, and in schools where you pack students close together and then try to have them all use their computers at the same time.

Is this a fundamental limitation of wireless? While it is true that there are some issues that cannot be solved, there are a lot of things that the network administrator can do to make the network work better. The key issue is the obvious, but under-appreciated fact that wireless networking is radio communications first. If your radio link doesn't work well, you have no chance of fixing it with your configuration and software. This paper is intended to give you an appreciation of what the issues are, and enough information to know what sorts of things to look out for when planning a high density wireless network.

Introduction and Overview

Wireless networks for conferences and schools tend to work very well when tested, and then collapse completely when all the users show up to use them. This pattern is repeated time and time again to the point where people tend to think that it's a fundamental limitation of Wi-Fi technology. There are real limitations that you have to deal with, but if you keep them in mind it is very possible to build a wireless network for thousands of people and have it be rock solid and reliable.

I have been running the Wireless network for the Southern California Linux Expo (SCaLE) since 2010 and this paper is based on the results of the attempts to provide wireless service from 2007 through 2012 at SCaLE. In 2012 we had 1965 attendees with 1935 unique MAC addresses on the network and 875 devices connected at peak.

The key thing to recognize when building a wireless network is that the network is primarily radios, and only secondarily digital. This doesn't mean that getting the radio side of things right will guarantee that your network will work, but it does mean that getting it wrong will guarantee that your network will not work.

Background

I

Prior to 2007, SCaLE only offered network capability to the booths on the trade show floor. To deal with rogue DHCP servers on the network, each booth is put in its own VLAN and network.

In 2007 SCaLE decided to start offering wireless network access. SCaLE purchased 10 3Com dual band access points and distributed them around the hotel. They discovered that these APs only allowed 32 devices per AP and as a result, they quickly became saturated and the wireless network was completely unusable.

In 2008, SCaLE borrowed 11 Linksys access points and the network administrator found a firmware hack that allowed him to boost the power. He put each access point on a different channel. This was also the year that the One Laptop Per Child laptops became available to the public and I volunteered at the OLPC booth demoing the laptops. The wireless network was unusable, in large part due to the interference between adjacent APs. After the event I spoke with the organizers and cringed when I heard how the wireless had been setup. I volunteered to help in the future.

In 2009, Xirrus donated wireless services to SCaLE. I don't have any details of what they setup, but the result was unusable.

In 2010, I received a call one month before the event, telling me that the commercial vendor that they had lined up to provide the wireless services had backed out, and since I had offered my expertise in 2008, was I interested in trying to run wireless services. I said yes and analyzed the problems.

I focused first on the radio side of the environment. If the radio side doesn't work, you have no chance of making the digital side work well, and you frequently won't even be able to tell what's wrong by looking at the network information.

Defining the Problem.

Radio Issues:

The 2.4GHz band (b/g) has 11 channels assigned in the US, but they overlap and as a result, you can only use 3 of the channels at once without problems.¹

¹ If you can use channels 13 and 14 in your county you can squeeze in a forth channel.

The 2.4GHz band is also used extensively by other equipment, including cordless phones, cordless microphones, bluetooth, and even microwave ovens. While the 802.11 protocol is designed to be resistant to interference from these things, these things can cause packets to be corrupted and result in retries.

The 5GHz band does not have these problems. Its channels do not overlap, there are far more of them (at least 10 available, sometimes more), and there are far fewer sources of other interference. Unfortunately, most equipment doesn't support the 5GHz band, and even when equipment does support it, many systems default to the 2.4GHz band.

Standard Wifi, like many mobile radio services, suffer from the Hidden Transmitters problem. A simplified description of this is where you have three stations along a line. The station in the middle can hear stations on each side, but the stations on the outside cannot hear each other. This prevents the stations on each side from avoiding transmitting when the one on the other side is already transmitting. When both sides transmit at the same time, the receiving station in the middle gets confused and can't make out either signal, causing both to have to retransmit the packet.

Excessive power levels can add to the Hidden Transmitter problem. It is common to think that if you can't get through, turn up the power, but if only one side turns up the power it seldom improves communications. This is because wireless networks are two-way conversations and if only one side gets louder it doesn't increase the range that the conversation can take place, but the stronger signal does go further and interferes with other stations.

The Wi-Fi protocols have evolved over time, with new modes being created that squeeze more data into a given amount of airtime. In most cases the newer, higher speed modes are more sensitive to interference, so the protocol includes fall-backs to slower modes when the data is not getting through. If the problem is outside interference, weak signal and similar problems, this works very well, but if the problem is an overload of the available airtime, the result is that each station transmitting takes longer to send its signal, which makes it more likely that a hidden transmitter or other interference will corrupt the packet resulting in retries.

802.11 has a fair amount of housekeeping traffic to let all stations in the area know that they exist and to maintain the connection to the Access Point. This traffic eats away at the time available and is frequently required by the spec to be transmitted at the lowest supported speed.²

802.11n can be a benefit or a problem. The fact that it can transmit more data in a given amount of airtime can reduce congestion, but if you enable the high bandwidth (dual channel) mode it will require

² Any broadcast traffic (such as SSID broadcasts, connection requests, etc) must be transmitted at the lowest speed supported so that devices that only support that speed and not higher ones will still be able to decode the message.

that two adjacent channels be allocated to it. Also, if the equipment is configured to operate in pure 802.11n mode, the 802.11b/g equipment will not recognize that there is a station transmitting and so will go ahead and attempt to transmit their packet.

Inappropriate use of high-gain antennas can be a problem as well. Unlike turning up the transmitter power, improved antennas help both transmitting and receiving the signal. But if they are used incorrectly they will cause the station using them to cover a larger area and so interfere with and be interfered by more stations.

Inappropriate access point/antenna positioning can have very similar effects to using the wrong antennas. It's very tempting to try and get the best possible coverage from each access point, but when you are trying to get the most number of users in a small area, this can actually hurt you. It's sometimes as simple as changing the height of the access point to limit how far it's signals will travel.

Mesh Networks require that the packets be transmitted over the radio more times, and as a result are almost always the wrong thing to use in a high-density environment.³

Retries can also be caused by problems on the digital side of things.

The Bufferbloat phenomenon⁴ where the delays in getting packets to their destination can result in the packets timing out before they arrive can also result in packets being retransmitted.

The typical collapse of wireless networks results from the combination of:

- Retries (frequently due to hidden transmitters or other interference)
- Fall-back to slow speeds
- Wasted packets (due to bufferbloat and other problems)

3 In this case I am referring to wireless links between the Access Points, In this case the traffic from many users is combined onto the uplink channel, the APs are using high power on the uplink channel and therefore the uplink channel is even more congested than the normal channels, resulting in them collapsing before the normal traffic does. Full 802.11s/OLPC style mesh networks collapse even faster as the packets are retransmitted over the normal channels.

4 In an attempt to prevent packet loss, and with memory becoming vastly cheaper over time, buffers on network devices have become very large. If there is significant congestion and the buffers stay full for an extended time, packets can sit in the buffers long enough that by the time they arrive at their destination they have already timed out and a replacement is in flight.

Solutions:

I needed to find out what I was up against. I did a site survey to find out what the situation was.

- Where are the network and power jacks (I've had cases where they were >8 feet apart)
- What other Wi-Fi signals are in the area, what channels are they on?

Some places in the Hotel were covered by 5 different existing networks, including networks from other nearby hotels.

Good tools to use are Wi-Fi analyzer on Android or Kismet on a laptop

- What interference is there in the area (usually not as critical as looking for Wi-Fi signals)

My-Spy spectrum analyzer can see all signal, not just WiFi signals.

- What effect do the walls have on your signal (movable partitions tend to block the signal more than traditional walls due to the metal mesh in the partitions)

I took an AP to plug in and then walked around nearby rooms and hallways to find out where I could hear it.

Once I knew what the environment is like, I worked to get as many access points in the area as you can get without them interfering with each other and without creating additional hidden transmitter situations.

The fundamental approach to making a lot of people able to use wireless in a small area is to use many low-power access points instead of a small number of high power access point. Cell phone service has a similar problem and solution, they refer to this as setting up microcells.

I encouraged the use of 5GHz channels. There are far more of them so you can have more radios covering a given amount of floorspace without interference, resulting in significantly more bandwidth per user. In addition to this making things better for the people who move to 5GHz, it also reduces the load on the 2.4GHz band, helping the people who don't move.

I turned power down on 2.4Ghz to allow for more access points without overlapping footprints.

I positioned the APs to take advantage of things that block the signal for me.

- The human body is mostly water, water absorbs 2.4GHz signals. By putting access points low in the room, the crowds will prevent their signal from going as far as they normally would.
- In the site survey I found out which walls block the signal. I was able to position some access

points closer to each other than I normally would have allowed.

I used advanced antennas carefully.

- I used directional antennas to cover a long theater room from the back of the room where I didn't have the ability to position access points more centrally.
- I also used mild directional antennas to direct the signal away from areas that were covered by other access points.

Digital Issues:

For SSID selection, I opted to use one SSID for each band (scale24 for 2.4GHz and scale5 for 5GHz), and re-use the same SSID on every access point. While putting a different SSID on each access point gives the user more control, the fact that using the same SSID everywhere allows the client to roam between access points as they move, without the user needing to do anything.




Illustration 1: LAX Westin 1st floor 2010

To make this work I configured the access points to act as bridges instead of routers, and run DHCP on a central server instead of on each access point. This makes it so that the IP address that a device gets continues to be valid as they move around the building.

Implementation.

2010

We purchased 16 5GHz NetGear APs (\$50 each) and 12 2.4GHz Fry's Electronics APs (\$30 each), along with three Cisco 160 APs that were purchased early on before testing the Fry's APs.




Illustration 2: LAX Westin 2nd floor 2010

We implemented normal site bandwidth saving tools:

- HTTP caching proxy (squid)
- Block streaming sources (DNS redirects to a placeholder page)
- QoS traffic shaping to allocate bandwidth between users ('users' being wireless users vs registration vs keynote streaming video, etc.)

Due to the extremely short timeframe, I turned down the power to 'low' with the stock firmware and crossed my fingers. I used directional antennas on the Cisco APs to direct their coverage area to minimize overlap with other APs.

Wireless worked well enough to crush the available Internet bandwidth (4.5Mb) for the first time. Overload caused the wireless network to be unusable.




Illustration 3: LAX Hilton 1st floor 2011

2011

SCaLE moved to new Hotel (LAX Hilton) with approximately 50% larger area to cover.

This hotel has 45Mb Internet connection. Prior to this SCALE conference they had never had it turned up above 20Mb. We had them enable full bandwidth, and kept it as close to saturated as our QoS settings would manage for most of the show.

I attempted to use just the equipment purchased in 2010, but with DD-wrt on the 2.4GHz APs. I was very unhappy with DD-WRT. It's more flexible and powerful than the stock firmware, but it does most of its configuration with special variables stored in NVRAM rather than with traditional *nix config files.




Illustration 4: LAX Hilton 2nd floor 2011

We had many problems with the 5GHz NetGear access points freezing, and with all the access points getting turned off, and with the network cables getting disconnected. The cables getting disconnected turned out to be the most disruptive problem, because the access points continue to advertise the SSID, but any clients connecting to them were dead in the water.

We found that we had insufficient access points to provide good coverage. When you were in a good spot things would work, but far too many areas and rooms did not work.

2012

SCaLE purchased 30 dual-band WNDR3700 APs (\$130 each). We ran them all on OpenWRT, with a custom compile. I kept the same locations for the access points on the first floor, and greatly increased the number deployed on the second floor.

I enabled Wireless Isolation. This prevents the wireless devices from talking directly to each other. This will break some use cases, but for the normal case where devices are talking to servers on the wired network it can both add some protection for the clients, and greatly reduce the wireless bandwidth needed. IP level broadcasts result in *many* retransmissions on wireless networks.

I lengthened the Beacon interval, it reduces the amount of housekeeping traffic, at the cost of it taking longer for devices to learn that the network is there or notice new APs as the users move around the building. With lots of access points there is enough overlap to minimize problems, and people are usually not moving that fast when heavily using the network.

I disabled connection tracking in the custom kernel. Connection tracking can be a very significant overhead on the CPU and RAM of the AP. Connection tracking is needed to implement Stateful Packet Filtering, but if you are not using any stateful firewall rules, it can




Illustration 5: LAX Hilton 2nd floor 2012

result is a significant amount of memory and CPU overhead for no benefit.

I set shorter than normal inactivity timers so that the APs don't spend resources trying to track devices that have moved or been turned off.

I adjusted kernel network buffers to be much smaller than normal to fight bufferbloat problems (50 instead of 1000). The Linux wireless stack includes quite a bit of buffering inside it, so setting the kernel buffers for the wireless interfaces very low helps minimize the possibility of excessive latency. There is some recent work in this area, but it does not yet deal with the buffers inside the wireless stack.

Running a web proxy also significantly helps fight the bufferbloat problem because by splitting the connections, you avoid having TCP connections with both high latency (long ping times of International Internet connections) and widely varying effective bandwidth of wireless connections. The widely variable bandwidth doesn't hurt much if it's only connecting to a local proxy server, and that proxy server has stable bandwidth out to the Internet. This is undermined to some extent by the "https everywhere" movement because https connections cannot be proxied.

As a result of the problems we had in 2011 with devices getting unplugged, we setup Nagios, cacti, and some custom rrdtool scripts to track what was happening on the devices.

We moved from having separate switch ports for each wireless band to having one connection per access point and using VLANs to isolate the administration, 2.4GHz, and 5GHz networks from each other.

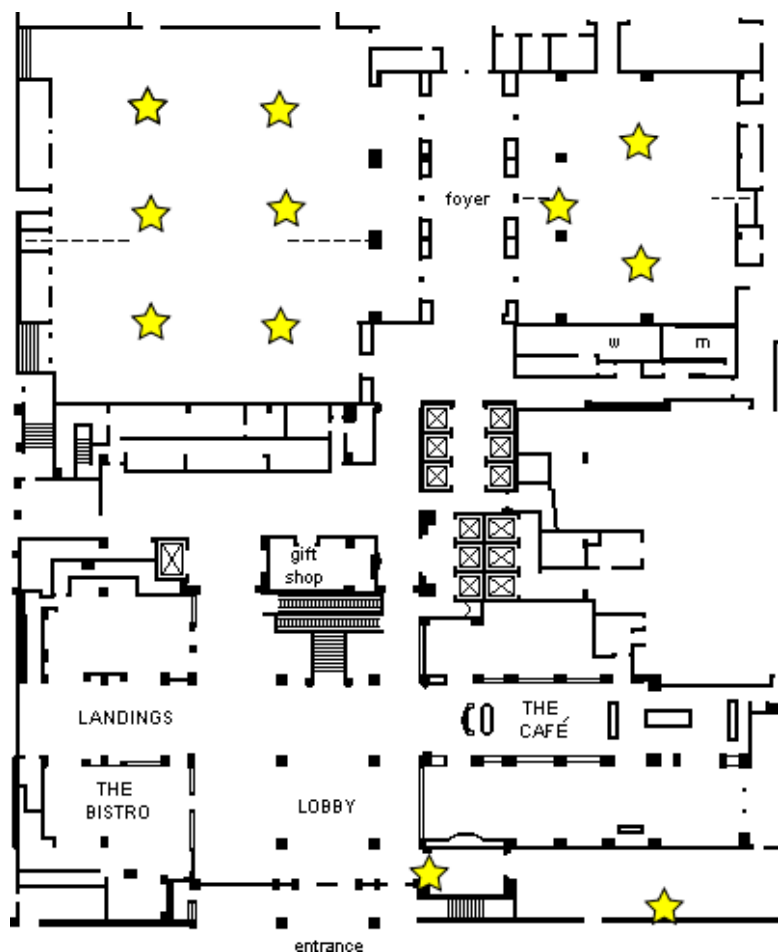


Illustration 6: LAX Hilton 1st floor 2013 (planned)

Results:

The wireless network was rock solid. Approximately 20% of the devices used the 5GHz band. Except for on the show floor, we had fewer than 70 devices connected to any one access point at one time. We had 1965 attendees over the three days with 1935 unique MAC addresses on the network and 875 devices connected at peak.

Future Plans

In 2013 we are expanding again. We will be taking over the rest of the hotel, expanding into more of the smaller rooms, using them for additional talks, BoF sessions, and streaming video from the main rooms to handle overflow. This will require an additional 20 access points, and the access points are going to be serving as the main

switches in some of the rooms, connecting the A/V gear to the network for the streaming video. We have not yet done the site survey for these additional rooms as of the time of writing, so we may end up turning off some of the radios on the access points in these smaller rooms.




Illustration 7: LAX Hilton 2nd floor 2013 (planned)

I will be looking to disable slow speeds. If you can disable the 802.11b speeds entirely you avoid having systems falling back to the extremely slow speeds and using more air time to transmit the same data, making the congestion problem worse. There are very few devices today that don't support at least 802.11g.

About the Author.

David Lang is a Staff IT Engineer at Intuit, where he has spent over a decade working in Security Department for the Banking division. He was introduced to Linux in 1993 and has been making his living with Linux (and using it as his desktop) since 1996. He is a Extra Class Amateur Radio Operator and served on the Civil Air Patrol California Wing Communications staff, where his duties included managing the state-wide digital wireless network. He has been running the wireless network at the Southern California Linux Expo since 2010. He is also active on various Open Source mailing lists. He can be contacted via e-mail at david@lang.hm, by phone at +1 818 292 7015

This paper and related materials are available at http://talks.lang.hm/events/LISA_2012/wireless