Whitepaper: The Bad Science of WLAN Design

A Whitepaper by Jon Foster

November 2014

Abstract

The current model for specifying, designing and installing 802.11 based Wireless networks (Wi-Fi) is fundamentally broken. Too little technical knowledge leads to a prevalence of ‘Bad Science’ style explanations, misused technical jargon and poorly designed networks. Ultimately, networks designed to these poor standards fail to deliver value and offer a sub-standard customer experience. The recent mass adoption and consumerisation of Wi-Fi means that this is no longer acceptable.

This whitepaper aims to demonstrate the reasons for this and suggests a way forward.

Introduction

802.11 Wi-Fi technology is enormously complex. Almost without question it is more complex than the more mature Ethernet LAN (Local Area Network) technology upon which it sits. One of the reasons for this is that, when designing a Wi-Fi network it is not possible to ignore medium in which the technology operates – electromagnetic radiation broadcast and received by radio transmitters.

On top of this, unlike with modern switched Ethernet networks, Wi-Fi transmissions must utilise complicated mechanisms to ensure that only a single radio is transmitting on a channel at any one time. The failure of these mechanisms leads to corrupt data and re-transmissions which serve to reduce the capacity of the network.

The standard industry approach to marketing 802.11 based wireless networks (WLAN) is to utilise rudimentary automated tools to suggest the locations for these radios (Access Points; APs). Once installed this wireless network is then automatically managed by network control software which dictates both radio channel and power settings.

The use of management software is a key feature of the normal approach to selling wireless networks, the theory being that its use will allow management of the network with little or no training or understanding of how the technology works.

Unfortunately, this almost inevitably means that the value of the network is never realised as the network severely underperforms.

The automated methods of network and radio management are not themselves cause of the problem - but are just symptomatic the wider issue - the whole "you don't need to understand it - it'll look after itself" paradigm. Under this paradigm, with customers having little advice available on how to plan and design networks, the radio management is left to find the *least bad* solution based on a poor physical layer design; the closest fit for the square block into the round hole.

Clearly this is attractive as it appears to provide the customer an up-front cost saving as it is claimed not to require expensive professionals to design, install and run the system. The reality is that the network which is delivered will not perform as expected and may require expensive remedial work in order to satisfy customer requirements.

This situation is likely to be exaggerated because poor levels of Wi-Fi technical knowledge possessed by the designers themselves means that the troubleshooting process is likely to be long. The delay between complaints from system users starting to appear and any effective remedial action being taken is likely to be significant. By the time that issues are resolved customers will already have moved on; not willing to accept a sub-standard service. Reputational damage will have been done.

Specific Problems

Bad Science and marketing

The result of this lack of industry knowledge is the prevalence of 'Bad Science' (a term popularised by Ben Goldacre[1]). Although somewhat inevitable with marketing material which must appeal to budget holders and less technically expert people, the lack of specialised knowledge within the Wi-Fi industry goes much deeper. Without a detailed understanding of how the technology works, network administrators and customers are forced to take marketing claims at face value – whether or not the benefits will be realised or even applicable in any particular situation.

By far the most common example of this is the use of headline "Data Rates" to sell technology uplifts. 802.11ac is touted as "gigabit Wi-Fi", and with the addition of "*MU-MIMO",* 802.11ac Wi-Fi is cited as similar to the move to switched Ethernet networks [2]. In other examples, vendors advertise fundamental parts of the 802.11 standard as *features*. These examples are discussed in more detail below:-

Data Rates

Whilst it is true that 802.11ac could provide throughput of multiple-gigabits, the achievement of these data rates pre-suppose an enormous set of features and environmental variables which are unlikely to prevail in the enterprise (such as 160MHz channels, a quiet and uncontended WLAN). A full discussion of expected performance is beyond the scope of this document, but it is more likely that we'd be looking at a real-world rate increase of 25% on 802.11n rates for a comparable scenario. For a signal stream device such as a smartphone this might mean achieving about 40Mbps in practice; far from true gigabit performance!

MU-MIMO for switched like performance

The 802.11ac standard describes methods by which a highly specified wireless transmitter can send data to a number of clients at once instead of one at a time. These methods are called "Multi-User Multiple In, Multiple Out" or "MU-MIMO". This has quickly been picked up and used to suggest that it will "transforms your AP from a hub into a switch for switched like performance". At the time of writing there are no devices on the market that support MU-MIMO, however from the standard we can see clearly that this transition to 'switched like performance' is very far from reality. Firstly, *clients* do not support MU-MIMO so therefore traffic **to** the AP will still be half-duplex, following standard Wi-Fi protocols. Secondly, the AP will still have a very limited number of radios to service multiple clients. A three radio AP would only be able to talk to three single radio clients at once, or a single two radio client plus a single one radio client.

When we compare that to a modern enterprise grade Ethernet switch which boasts "24 ports of non-blocking performance" [3], we can see that the claims of "switched like performance" are, in fact, very misleading.

"Clear Channel Assessment to avoid channel conflicts"

One vendor states, of their equipment [4]:

> *Clear Channel Assessment (CCA) avoids channel conflicts using its clear channel selection feature and fully realises the advantages of channel binding, greatly enhanced the wireless performance.*

When we try and analyse what this means we find it is quite a mix up of real features of the 802.11 standard - Clear Channel Assessment, and a mashup of technical sounding terms the meaning of which is impossible to determine.

CCA is actually a mechanism which forms part of the 802.11 standard; it's fundamental to the operation of *every* piece of Wi-Fi hardware. Touting it as some kind of special feature is a little like suggesting that office space "comes with oxygen for enhanced productivity".

The rest of the terms are vague. What are channel conflicts? Is the clear channel selection feature the same as the CCA, or something different? What is channel binding – does this relate to the ability to configure wider 40MHz channels with the 802.11n standard? If so, again, the CCA function is just functioning as a basic part of the standard; not some special value-add feature.

It would be naive to think that large corporations are not going to use very expensive and very clever marketing strategies to highlight any possible advantage that may be gleaned through using their products. The issues only really come when purchasing decisions need to be made. Somebody is required the ability to evaluate or judge whether the claimed benefits will apply to them. When purchasing shampoo it isn't really that important that the customer understands the *bendy-holdy technology* [5] but when purchasing a multi-million pound WLAN, arguably, it is.

Vendor Training

Typically, vendor training is purely geared towards control, configuration and troubleshooting of the equipment itself and not to the fundamental principles of 8021.11 WLAN operation. A wireless solution will be described and sold as a system which, once installed, effectively manages itself. Neither the sales team nor the customer are likely to be any the wiser.

Given this, if any discussion of network capacity is required, then it is likely to result in the over provision of APs. And, given the lack of awareness of the problems this causes the sales team can do this with a clear conscience - using what they believe to be "tried and trusted" accepted practices (often, ironically, described as "best practice").

Another problem created by this focus on Vendor training is that this leaves engineers with very large gaps in their knowledge relating to troubleshooting. Troubleshooting end user connection issues often requires a site visit and analysis of information only available in-situ – and completely un-related to the actual vendor's equipment. Vendor training very rarely covers this in enough detail to be useful. This shortfall will often lead to expensive mistakes through misdiagnosing the cause of problems and delays in their resolution.

Quick Fixes – deploying another AP

We encounter yet more issues when we consider the most frequently quoted solution to perceived capacity problems (i.e., a slow WLAN) - deploying an extra AP. This is best illustrated by looking at AP 'beaconing' a fundamental element of the 802.11 standard.

By default each wireless network (SSID) broadcasts its presence and capabilities in a 'beacon' every 100 milliseconds, using up valuable network resources. An enterprise may run a "Staff" SSID a "Student" SSID and a "Guest" SSID for example. A WLAN Administrator will quite often be pressured by the business into broadcasting more SSIDs on their WLAN. Without a good level of training the administrator would not likely be aware of the impact of adding these extra networks – all the beacons from these networks will use up wireless network capacity, or "airtime". If multiple APs within "earshot", where one AP can detect the signal from another on the same channel (as if often the case with high capacity and over provisioned designs), the situation can develop where this beaconing alone can use up virtually *all* the available resources on the network. This would present as a lack of capacity to the untrained administrator. Given this lack of knowledge, this individual will often seek to resolve the issue by deploying another AP closer to user groups who are experiencing poor wireless service. This would add yet more beaconing overhead and actually making the situation worse!

It is also worth noting here being able to broadcast up to 16 separate SSIDs' is actually touted as a selling point and included on marketing material[6] – with no mention that it will cripple network performance!

Auto Planning Tools

A number of the largest equipment vendors include some rudimentary planning tools which are typically used to plan AP locations. The process normally involves loading a basic floor plan, selecting a few options and watching as the planning application quickly covers the whole area in pretty colours which appear to show a working WLAN with optimal coverage. This approach reinforces the view that this will be sufficient to design a wireless network; everything else will be taken care of by the system itself. These tools are simple enough to be used by a novice network administrator, but…

Early versions of these tools included elements such as "coverage overlap" of up to 200% with no technical explanation of what this actually means – 200% of what? Following this design strategy would have resulted in a very dense AP deployment which would have been expensive to purchase and is seriously technically flawed. Given that 90% of users (those on the 2.4GHz network) would share only three channels no matter how many APs were installed. Very little extra network capacity created by the extra equipment. Even if these AP's all reduced their power levels in an attempt to compensate (as below), it is more than likely that far more APs than necessary would be installed.

Automated RF Management

Given the limited number of non-overlapping channels available (especially at 2.4GHz) it is important to plan the allocation of channels to minimise interference caused by APs in close proximity to one another operating on the same channel. A number of the major vendors have attempted to allow for the planning of channels without the planner having any knowledge of the subject matter by integrating automated radio management software into their system controllers. The theory is that this allows their radio management software to ensure that channels and power levels are optimally configured.

Whilst it is true that this technology does let a network administrator deploy a network with absolutely no understanding of a 'channel' or required power level of an AP in 'mW' or 'dBm', it also hinders them from designing an optimum configuration. The network may well configure the APs (eventually) with a stable set of power and channel settings but often this is far from optimal. In some cases adjacent APs end up both using the same channel and thus mutually interfering. Whilst this may be the '*least bad'* design, it surely implies that the WLAN has been very poorly designed with respect to physical placement of APs – hardly delivering quality or best value. Indeed, it leads to the very misguided believe that you can install literally whatever you like and leave it up to the magic radio management to make it work.

Further, when utilising this technology, the job of troubleshooting a problem is often elongated significantly as these networks rarely stay in a stable configuration as settings are continuously dynamically changing.

Some vendor descriptions include suggestions that the network is “self-healing”. This technology refers to the system’s ability to increase (or decrease) the power level of an AP. If a device fails, other APs can increase their power levels to compensate. The reality is though that this power increase is “Omnidirectional”, that is, it transmits its signal in all directions. This is because the antennas which have been installed are almost always Omnidirectional (again for ease of deployment). If we’re really lucky we might find that an AP can increase its power to cover a hole created by a nearby AP failure, but in every other direction we have almost certainly just created large amounts of Co-Channel Interference (CCI) and quite possibly, a cascade of AP channel changes as the network tries to “Self-Heal”. In some cases this can cause as large a service degradation as the original AP failure. Given that vendors work hard to ensure the mean time between failures (MBTF) of their equipment is as long as possible and that failures are rare, does the network *really* have to ‘self-heal’?

Add to this the fact that in order to support good 802.11 ‘roaming’ we are likely to have a coverage overlap with another AP offering service to any particular area, would we not be better off following a similar paradigm to the LAN where an equipment failure triggers an alert which mobilises an engineering response, replacing the faulty kit in the next window of opportunity?

Arguments in favour of this type of technology often include statements such as “your busy administrators have not got the time to manually configure channels” and “our technology allows your wireless network to avoid interference from individuals deploying their own unauthorised APs.”

Let’s take a moment to consider these statements.

Network administrators have not got the time to manually configure channels

A significant amount of work goes into the deployment of a new wireless network – even if we apply lip-service to the design process, plenty of effort will go into cabling and documentation. Given this, is it really that much of an effort to manually configure a channel and power setting to an optimum? Surely a tiny fraction of the entire project effort. A small price to pay to *ensure* delivery of maximum value?

(Note that there are some scenarios where it makes sense to allow an AP to select the best available channel itself, for example a branch office or coffee shop environment with a very limited number of APs deployed and a high likelihood of competition from other WLANs in the locality.)

Avoiding interference from unauthorised APs

There are only 3 channels to choose from on a 2.4GHz wireless network. When we say “Avoiding interference from unauthorised APs”, what we’re really saying is: we are willing to sacrifice our optimum channel configuration in order to cater for a device which is at best competing, or at worst a rogue? There is no choice but to move to one of the other two available channels which are all but certainly already contested. One channel change can quickly lead to other APs changing channels in a cascade effect and further, on 2.4 GHz networks with only 3 channels to work with, there is rarely anywhere to really *go*.

Signal Strength Surveys

Another key result of the lack of true technical knowledge at the administrative, reseller and integrator level is the prevalence of the *Signal Strength Survey*. If nothing else is known to the WLAN administrator, they will likely have heard of Signal Strength. A short time on Google can tell you that Cisco suggest a signal strength of -67dBm for a voice capable network. In principle, armed with this bit of information one could commission a WLAN, positioning APs according to signal strength.

I would suggest that this is little better than the *"3 bars on my iPhone approach"*.

Whilst a 'strong signal' is one metric which assists us in planning a WLAN it is just as one of many. More APs per area would likely equate to a client receiving a stronger signal, but this says nothing about the quality of that signal; whether the network is likely to have enough capacity; whether the network is likely to contain problems such as 'hidden nodes'; whether there is interference from non-Wi-Fi sources or a myriad of other shortcomings.

Clearly then, *Signal Strength* is not enough to plan with.

A Hard Sell

All of the above issues are compounded by one further problem; attempting to sell an alternative methodology to customers which takes more time and has higher up-front costs (represented by increased professional services), in most cases is a very hard sell. The above issues are very technical in nature and with few exceptions, customers will not be well placed to differentiate well-argued counter case from empty sales patter.

Any project manager asked to justify the extra spend would likely be forced simply to pass on the word of the consultant to an even less technically able budget holder who is driven to deliver value first and foremost as a financial measurement.

So, a Broken Model?

The preceding sections have highlighted the issues which ensure that many Wi-Fi networks are sub optimal. The whole WLAN design and deployment ecosystem is geared to hide the standards based knowledge and concentrate on the vendor specific technology which attempts to automate everything – and do away with the need for an in depth knowledge of the technology.

The rapid adoption of Wi-Fi and market penetration of Wi-Fi enabled devices is increasingly meaning that sub-standard designs are no longer adequate. At one end of the scale, as soon as a poorly designed network is used in any meaningful way problems develop which are expensive to solve. At the other end of the scale, those networks that never see high levels of utilisation may simply be over-specified with large numbers of superfluous AP's. The total cost of owning these devices includes not just the AP itself, but the cost of deploying structured cabling, software licencing, not to mention the LAN switch ports and power required to keep the extra units running.

Clearly this situation has not come about by deliberate design; equipment vendors have needed to create solutions to meet rapidly growing demands. But how does the industry move forward to ensure that the networks it delivers to customers are well designed, providing the optimum level of capacity and value?

The Solution

On its own, technology cannot fill this gap; the radio frequency management algorithms are evidence of this; going some way to make the best of a bad situation, but are in themselves, inadequate to eradicate the problems caused by poor physical layer design.

The problem is not, however insurmountable and can be addressed in part with adequate training and the use of knowledgeable and experienced engineers during design and implementation. Following the training models used for LAN engineers is not enough; a detailed knowledge of the *physical layer* is required, alongside detailed knowledge of the 802.11 standards - even to guarantee reliable delivery of basic services (Web/e-mail, etc.) in a high capacity WLAN environment.

The CWNP

The first part of the solution to this problem is ensuring that technical staff and stakeholders within the sales channel and specialist Wi-Fi consultancies and enterprises, where possible, have an adequate understanding of the 802.11 standards alongside their vendor specific accreditations. For this a neutral framework of certifications such as the CWNP[7] framework is required. This framework ensures that its certificate holders possess a good understanding of all wireless operations from the physics of RF communications upwards, both theoretically and in practice – and not just the specifics of operating any particular vendor's equipment. Given this knowledge, many of the errors outlined above would not be made. A WLAN professional armed with this knowledge has many of the tools required to avoid having to rely upon automatic radio management systems and their often desperate struggles to find the least bad solution to fundamentally flawed network designs.

Maintaining highly skilled, specialised IT professionals is not an option for everyone. Even those enterprises large enough to employ a network specialist are unlikely to ensure the individual is skilled enough to cover all base to and adequate depth. In these situations a trusted partner can be used to provide missing skills. In these cases the CWNP framework is even more valuable as they will highlight levels

The Future

At least one vendor, Aruba Networks, has made some small steps in the right direction and has at least acknowledged the existence of the CWNP [8], but there is still a long way to go. Without a fundamental change of their ethos, buy in and training internally there is still likely to be an over-reliance on the "auto-configuration paradigm".

Unfortunately, the CWNP framework on its own does not go the whole way to ensure that consumers of professional services are protected from jargon. Although the CWNP set the standard for vendor independent knowledge of subject matter and their certifications do require real world experience to obtain, they do not specify at all *how* one delivers a wireless network or, for example, what should be included in a technical audit of an existing one. Therefore, arguably, another element is required; a standardised framework of best practice. Such a framework to ensure that the certified Wi-Fi practitioner is carrying out an internationally agreed set of practices, tests and measures to ensure quality networks and meaningful results.

How the industry tackles these outstanding issues is beyond the scope of this paper, but the requirement for change is growing. Wi-Fi is no longer a "nice to have" or "value add" technology but has moved into the core of the Enterprise. With widespread adoption of ‘Mobility’, ‘BYOD’ and deployment of applications such as Microsoft Lync a fully functional and highly performing WLAN infrastructure will be essential.

Credits and references

[1] http://www.badscience.net/

[2] http://arstechnica.com/information-technology/2014/05/wi-fi-networks-are-wasting-a-gigabit-but-multi-user-beamforming-will-save-the-day/2/

[3] http://encyclopedia2.thefreedictionary.com/non-blocking

[4] http://www.tp-link.com/lk/products/details/?model=TL-WN722NC

[5] A tribute to a UK advertisement from the 1990's which at least acknowledged that the typical pseudo-science in adverts has little value. https://www.youtube.com/watch?v=hUkUjdaSe90

[6] http://www.ruckuswireless.com/products/zoneflex-outdoor/zoneflex-7762

[7] http://www.cwnp.com/

[8] http://community.arubanetworks.com/t5/Training-Certification-Career/Becoming-Mobility-Certified-CWNA-June-2014/m-p/160672/highlight/true#M911

About the Author

This white paper was written by Jon Foster a Wireless LAN administrator working for the University of Exeter in the UK. The views expressed here in no way represent those of the University of Exeter.

E-mail: jon@weaponsgradewifi.com

Web: http://www.weaponsgradewifi.com/