


The Overlapping Channel Problem

What Overlapping Channel Problem?

In the UK, the area of the wireless spectrum set aside for the use of 802.11b/g wireless networking devices is the ISM (Industrial, Scientific and Medical) band between 2.400 GHz and 2.497 GHz. In the UK this is subdivided into 13 channels of 25 MHz. In the US, only the first 11 of these channels are available – a fact with implications for UK deployments (see ‘WAG’s Advice’ below).

However, this description can give a misleading picture. Radio signals are not tightly confined to a single point in the spectrum but are distributed symmetrically around a ‘mid’ frequency. This will be familiar from the experience of tuning a radio, where the signal gradually gets stronger as you turn the dial before hitting the peak frequency and then slowly dies away again. In fact, the width of spectrum across which the signal is distributed is related to the bandwidth (in the networking sense) of the data being transmitted and the symbol rate used to encode it (this relationship is discussed in Nyquist/Shannon sampling theory). 802.11b, with a bandwidth of 11Mbit/s and a binary symbol system, requires 22MHz per channel (802.11g achieves its greater data rate in the same channel width by increasing the bits carried per symbol). So, a better picture of the channels in the 2.4 GHz spectrum is given by the following (the four differently shaded, solid curves indicate signals associated with the four non-overlapping channels: 1, 5, 9 and 13):

Channel distribution in the 2.4GHz ISM Band



Channel	Start Frequency	Mid Frequency	End Frequency
1	2399.5 MHz	2412 MHz	2424.5 MHz
5	2419.5 MHz	2432 MHz	2444.5 MHz
9	2439.5 MHz	2452 MHz	2464.5 MHz
13	2459.5 MHz	2472 MHz	2484.5 MHz


Significant overlap can be seen between adjacent bands. APs (access points) with physically overlapping broadcast cells configured to the same or near-adjacent channels will be trying to use substantially the same frequencies to transmit and receive data.

Whilst a problem, overlap does not completely preclude wireless communications. The Media Access Control (MAC) specification for 802.11b networks implements a CSMA/CA (Carrier Sense Multiple Access Collision Avoidance) mechanism, which effectively means that each AP listens on its channel before attempting a transmission. In the overlap scenario described above, the multiple APs would end up sharing the available channel resource (with some further reduction in throughput due to collisions). WiFi was *designed* to support many interfering, overlapping networks, handling the resultant packet collisions

with minimum fuss. So rather than entirely disrupting a wireless LAN, overlapping channel assignments ‘merely’ greatly reduce its efficiency. This is clearly not desirable, and so the standard advice in wireless networking circles in the UK is to select one of the four non-overlapping channels for adjacent cells, allowing all APs to operate at their peak throughput. However, this raises a new problem ...

The Four Colour Problem

The ‘Four Colour Problem’ is a mathematical curiosity dating back to 1852, when Francis Guthrie, trying to colour in a map of counties of England, noticed that four colours sufficed to ensure that adjacent regions (i.e. those sharing a common boundary segment, not just a point) received different colours. However, a rigorous mathematical proof for this conjecture was not forthcoming until that by Appel and Haken in 1976,¹ itself controversial for its use of computers in proof.



It is not a great leap to substitute the pattern of adjacent cells of a wireless LAN in a building for the counties map, and the available frequency bands for the colourist’s paints. So, mathematics seems to be telling us that whatever the pattern of adjacent wireless cells on a plane surface, they can always be configured in a way that completely avoids channel interference, using only the four non-overlapping channels and avoiding adjacent cells using the same channel. (In fact, in typical situations, even three channels can be enough – fortunately for the Americans, who only have 11 channels in their corresponding ISM band and thus only three non-overlapping channels.)

Unfortunately, we seldom deploy wireless on just a single floor of a building, and adding the third dimension into the problem makes it much harder to find a non-overlapping solution. In most buildings, wireless networking signals penetrate ceilings and floors as readily as they do walls (most omnidirectional antennae on APs are $\frac{1}{4}$ wave whips, which generate a flattened torus of signal distribution centred on their antenna; PCMICA client radios tend to emit in a spherical pattern), so our ‘colouring’ problem channel configuration design has to address cell boundaries both above and below as well as on the same level as the AP.

Somehow we need to find more ‘colours’ for our palette.

A Solution?

An analysis by Mitch Burton² reminds us that whilst using more closely spaced channels does result in interference, the actual loss in throughput is minimal if we select five widely spaced channels instead of the usual four (or four rather than three, in the States). According to Burton’s analysis, when two ‘buffer’ channels separate each of the 802.11b channels in use (i.e. channels 1,

4, 8, 11 and 13), any given AP will only experience about 4% more interference. In comparison, the normally-recommended three channel buffer (i.e. using 1, 5, 9, 13) results in <1% interference between APs. Taking into account the fact that the current generation of wireless equipment also uses filters to further reduce the impact of interference issues, it seems safe to conclude that the overhead incurred through using five channels rather than four of the available thirteen is acceptable.

Another Solution?

Most enterprise-class APs offer the capacity to assign an appropriate channel automatically. This is typically determined during the boot process by a simple poll of each channel at a randomised interval, so that groups of APs booting together do not all select the same band. The channel assignment from these polling results is made by a very basic algorithm, picking either the channel farthest from all those already in use, or if all are used then picking at random.

Some manufacturers also offer wireless management devices that control a group of APs and, among other functions, make channel and transmission power assignments taking into account the group as a whole. These systems can also monitor the spectrum and dynamically adjust the channel assignments in response to changes such as new interference sources or AP failures (e.g. by extending range to compensate as much as possible).

WAG's Advice

Firstly, the JANET WAG (Wireless Access Group) would advise against using automatic channel assignment options, either at the AP level or as implemented by a wireless management device. For the former, the primitive algorithm is barely sufficient to avoid initial interference on booting, and is otherwise little better than a guess. Management devices make a better stab at designing a plan, but tend to be biased towards (a) US channel assignments and (b) using reduced output power to create a viable pattern. This results in a greater number of APs being required to cover an area than might be required with manually-assigned spectrum management. Most importantly, a manually-assigned deployment is understood by its operators, whilst automated systems encourage treating spectrum management as a 'black box'. A known wireless environment facilitates planned upgrades, sensible monitoring strategies and reaction to new sources of interference. That said, the technology continues to evolve, and this area should be monitored as firmware upgrades become available in case a viable technological solution arises.

Returning to manual assignment, in most WiFi deployments in the UK it is sufficient to adopt the convention of using the four non-overlapping channels (in fact, in many situations, three channels are enough to achieve a non-overlapping topology, so an even greater buffer between adjacent cells can be achieved). However, there are scenarios when even four channels are insufficient:

- in higher density deployments
- across multiple floors of a building
- where external interference precludes the use of one or more of the 'usual' channels.

In these circumstances it may prove possible to use partially-overlapping channels at the price of only minimal overhead in terms of traffic collisions. The rule of thumb is to maximize the channel separation between cells as much as is practical, and to be aware that once the channel separation drops below around 25MHz (three channels-worth of buffer), you will start to experience degradation in performance. Unfortunately, wireless deployment is something of an arcane art, and each case must be evaluated on its merits and tested *in situ* to see whether the impact of the increased overlap is acceptable.

You should also bear in mind that reducing the power output of your APs is another way to decrease the number of interfering cell boundaries, and that when all else fails the 5GHz spectrum used by the 802.11a standard is available, with an increased number of non-overlapping channels.

In the situation where an additional AP is required within an existing deployment, the efficiency of the WLAN is best maintained by redesigning the entire channel assignment pattern for all APs in the area rather than taking the easier option of just optimising the band chosen for the AP being added in.

The potential difficulties of managing channel assignment in wireless deployment also make it advisable for institutions to develop a formal spectrum management plan. This is designed to:

- prevent rogue deployments occupying channels you need for future 'official' WiFi services
- develop relationships with neighbours who might also deploy wireless (alternatively, try to get your kit deployed first)
- check that unrelated technologies that also use the ISM band (security systems etc.) do not limit your future wireless networking options
- develop networking policies and monitoring regimes that let you police the radio environment on campus to guard your local spectrum resource.

Finally, as noted above, there are two more channels available for WiFi hardware in Europe than in the USA. Given the dominance of the US market for this technology, it is perhaps unsurprising that there is some observed bias in hardware performance around this point of difference: we observe that some hardware exhibits interoperability problems on channels 12 and 13 that it does not display on the lower, common channels 1-11. Whilst these difficulties often derive from a 'factory default' assumption of US deployment that is readily corrected through configuration change or driver updates, it is worth noting the potential for increased support load from deployments making use of these channels where a common hardware platform is not provided (and as a point to test when assessing hardware for deployment).

1. K. Appel, W. Haken, and J. Koch, Every planar map is four colorable. II. Reducibility. *Illinois J. Math.*, 21(3):491-567, 1977
2. Mitch Burton, Channel overlap calculations for 802.11b Networks. TechOnline, November 2002. http://www.techonline.com/community/ed_resource/tech_paper/21668