Frequency Servers

KEVIN MITCHELL
Agilent Laboratories

As WiFi grows in popularity, and public hotspots become more widely deployed, there will be increasing pressure to allocate additional bands and channels for 802.11-style networking. The paper outlines a framework for automatically controlling the frequencies, power levels and modulation schemes used by mobile communications devices such as 802.11 access points and stations. We propose the introduction of *frequency servers* that map geographic locations to the frequency bands, power levels, and modulation schemes a particular service is permitted to use at these location. A *beacon*, a low powered transmitter with a known and trusted location, uses the information provided by the frequency servers to control the spectrum usage of devices within range of the beacon. Cryptographic signing is used to authenticate both the information provided by the servers, and the location of the beacons.

Key Words and Phrases: Spectrum allocation, WiFi, 802.11, beacons, location servers, software-defined radios

1. INTRODUCTION

As WiFi grows in popularity, and public hotspots become more widely deployed, there will be increasing pressure to allocate additional bands and channels for 802.11-style networking. Furthermore, the eventual introduction of 3G phones is unlikely to quell this demand. The data rates obtainable with such phones are still relatively low, and so hybrid solutions will almost inevitably be deployed. These will use the high-speed 802.11 connection when in range of a hotspot, and the 3G network, at lower speed, elsewhere. The resulting contention within popular hotspots will, in the short-term, be resolved by adding additional access points and making the cell sizes smaller. There is a limit to how far you can go in this direction before some of the commercial advantages of WiFi are undermined. A longer term strategy for easing contention might be to allocate more channels to such services, with each access point supporting multiple bands.

One of the limiting factors in allocating more bands to WiFi services is the differing uses of the wireless spectrum across the world. This is easily illustrated by looking at the current 802.11b channel constraints. There are currently 14 channels defined, between 2.412 and 2.483 GHz. In North America channels 1 to 11 can be used. In Europe you can also use channels 12 and 13. Except that if you are in Spain you can only use channels 10 and 11, and in France 10, 11, 12 and 13. Finally, in Japan you can only use channel 14. Such constraints are particularly onerous when many users of WiFi are mobile. A user roaming between countries needs to be aware of the local constraints. When a device is used in infrastructure mode then the problem is less severe as the access points are typically not mobile, and

Author's address: K. Mitchell, Agilent Labs Scotland, South Queensferry, Scotland EH30 9TG

will be configured to respect the local constraints. When devices use ad-hoc mode to communicate then the problem becomes more acute.

Most countries have holes in their wireless spectrum, frequency bands that are either unused, or where the incumbent users could be easily moved. However, in many cases a hole in one country will be used for other purposes elsewhere. Coordinating the freeing up, and reclasification, of bands across countries is therefore a time-consuming process. Allocating additional bands on a per-country basis is obviously easier, but increases manufacturing costs, and confusion. The FCC's proposed spectrum shakeup [Kwerel and Williams 2002] will potentially make the situation much more complicated in the short-term. WiFi may be allowed to spread into many bands in the US that will conflict with other uses in the rest of the world. Roaming users may then cause severe problems when they take wireless devices intended for the US market overseas. An automatic mechanism for choosing channels based on geographic location may make such a strategy more persuasive.

Even within a country there are many potentially unnecessary restrictions. For example, within an office block it might be feasible to use the marine or ham radio bands at low power for 802.11-style communication, without causing interference to legitimate users. However, because there is currently no automatic mechanism for preventing use of such bands when a laptop leaves the building, the authorities naturally err on the side of caution. A mechanism to automatically constrain channel usage to within relatively fine-grained geographic areas could enable such usage.

In the past WiFi stations have been constrained to a small choice of bands, typically one. Even dual band cards, operating at 802.11a and 802.11b frequencies, are relatively rare at present. So the possibility of being allowed to use hundreds of channels in many different frequency bands has been largely academic. But current work in software defined radios, and devices such as SiRiFIC's wireless products[SiRiFIC 2003], may make it far more feasible to deploy such devices. This proposal outlines a mechanism for using such technologies in a safe fashion across many geopolitical areas. The aim is to ensure such radios only use an appropriate set of frequencies, power levels and modulation standards given their current location.

2. ADAPTIVE TECHNIQUES

One approach to exploiting spectrum holes is to use adaptive techniques (e.g. [Motorola 2002]). Radios listen for signals in a particular frequency band, and if no traffic is heard over some period of time, or the receiver detects the band is already being used for 802.11 communication, then the station takes this as implicit permission to use this band. Such techniques have the advantage of being very distributed, but they also have some obvious disadvantages. For example, consider the situation illustated in Figure 1, where station **A** is trying to detect potential spectrum holes. It is outside the range of station **C**, and so will be unaware that this station is transmitting data to station **B** on a particular frequency band. If **A** erroneously concludes that this frequency is unused, and starts transmitting on this frequency, then it may disrupt station **B**.

Peer to peer networking techniques can go some way towards alleviating such problems. If adjacent nodes exchanged their knowledge of potential spectrum holes

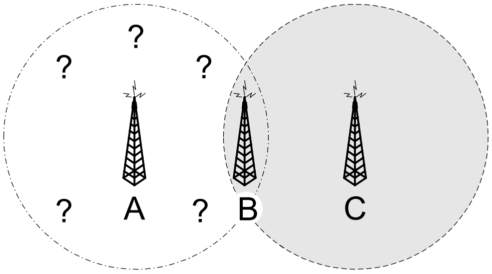

Fig. 1. Hidden nodes

then a node may be able to build a better picture of the geographic extent of such a hole. If the extent is substantially larger than the transmitted range of the nodes wishing to use the hole then this would give the nodes more confidence that there are no hidden node problems associated with its use. Of course it still gives us no guarantees, and bands that are used intermittently may still be misdiagnosed as holes. Furthermore, when an official user does start transmitting on the band there is no easy way to distinguish this use from other, unlicenced, users with no more rights or priority to use the band than the WiFi stations. Under what circumstances should these stations therefore stop using such a band when other transmissions are detected?

Adaptive techniques make it difficult to provide any guarantees over when and where a station will use a particular channel, and at what power levels. This makes it unlikely that governments would be comfortable with large-scale liberalisation of the spectrum if they had to rely on such technology. But what if we could provide such guarantees...

3. FREQUENCY SERVERS

The proposal outlined in this document describes a framework for automatically controlling the frequencies, power levels and modulation schemes used by mobile communications devices such as 802.11 access points and stations. It can be viewed at one level as a cross between DNS[Mockapetris 1987a; 1987b] and DHCP[Droms 1997]. Our first requirement is for governments adopting this approach to set up *frequency servers*, analogous to root domain name servers, providing the following service. Given the coordinates of any geographic location within the country the server is able to determine which frequency bands, power levels, and modulation schemes, a particular service is permitted to use at this location. Initially such servers might just treat the whole country in a uniform way, returning the existing 802.11b channel details for the WiFi service, for example. But over time the database would be extended with more liberal rules for different areas of the country, down to the granularity of individual office blocks for example. These extensions would be added primarily on a demand basis. Corporate users would apply to have their sites surveyed, for a fee. In return they would be able to use a wider range of channels, and hence bandwidth, within the geographic constraints of their site. Just as root domain-name servers delegate to other servers that cache some of the details, a hierarchy of frequency servers could also be constructed, with root servers just containing default settings for the whole country, and local servers providing more specialised extensions and constraints based on local knowledge.

How does a wireless station determine which channels can be used in this scheme? It would be tempting to use something like DHCP for this purpose. If a DHCP server knew its location it could obtain channel usage information from the frequency server. An additional DHCP option could be defined that provided this information when a wireless station retrieved an IP address. Unfortunately there is no guaranteed tight correlation between an IP subnet and its geographic location. To avoid such problems we propose the introduction of *frequency beacons*.

A beacon, in the context of this proposal, is a low-powered transmitter, attached to the Internet, with a known and trusted geographic location. Some beacons would have GPS receivers built into them to determine their position. However, in many cases beacons will be sited inside buildings where GPS is not available. An alternative approach in such situations would be to hardwire the position into the device. In both cases cryptographic signing would be used to ensure that the positional information could be trusted. For example, in the case of a hardwired position an inspector would certify that the position recorded was accurate, and download a digital certificate into the device to authenticate this information. The device would have a tamperproof mechanism that would invalidate the digital certificate if the device was moved.

The beacon would periodically interrogate the frequency server to determine the current frequency band policy for its vicinity. Such information would have a limited lifetime, just as with DHCP leases, ensuring that beacons do not cache this information indefinitely. The beacon and the frequency server would authenticate each other to ensure that the information returned by the server is genuine. The beacon would periodically transmit this information to all wireless devices under its control. The intention is that all access points and stations within range of the beacon could use any of the channels permitted by the beacon.

There is an obvious trade-off here. If a beacon transmits the channel usage information at high-power, i.e. its cell size covers a wide area, then there is a large risk of stations within this cell conflicting with other “official” users. The number of channels that can be used in this area will therefore have to be restricted to the “standard” set. A small cell size would cover a smaller geographic area, and so could use a wider choice of channels without causing interference to others. The drawback is that more beacons would need to be deployed. The frequency server uses the beacon’s cell size, as well as it’s location, when determining an appropriate set of channels and other policy data. Of course the area covered by a beacon will be a complicated shape, governed by many factors. To ensure we err on the side of safety the range of the beacon, for a given power level, is overestimated, and the coverage area assumed to circular, with a radius determined by this range.

If a beacon can transmit at varying power levels then logically it can be viewed as a set of beacons, each with their own power level and resulting cell size. The beacon may interrogate the server once for each cell size, and then broadcast the appropriate policy for each power level. Devices close to the beacon will receive multiple broadcasts, with differing policies. However, there will be a natural ordering between these, and the device will base its decisions on the most liberal policy

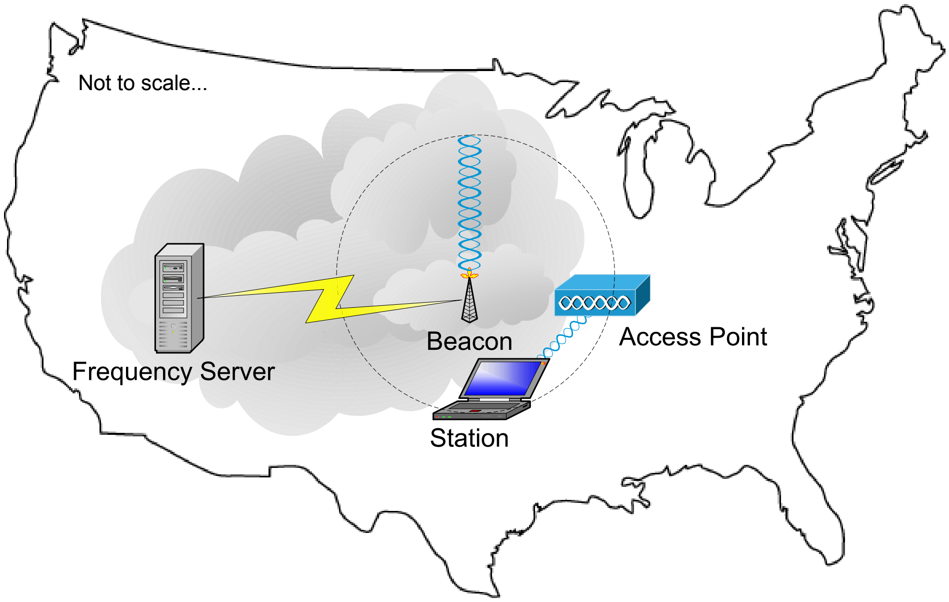

Fig. 2. Beacons

received.

The cell size requires further definition. The contract between beacon and server has the following form. Any device within the bounds of the beacon's cell, e.g. within a specified distance of the beacon, is permitted to use the specified wireless resources, as long as any signals escaping outside the cell are guaranteed to be below specified thresholds. Suppose the beacon always transmitted at maximum power, and the maximum distance at which this signal could be received was B_d. Consider a station at this distance from the beacon. If it used one of the channels permitted by the beacon, also at maximum power, then let's assume the maximum distance at which this signal was above the threshold was S_d. The maximum distance at which a device controlled by the beacon could disrupt other users would be $B_d + S_d$. This would therefore have to be the cell size reported to the frequency server. The situation is illustrated in Figure 3. There are some tradeoffs here. The larger the cell size the smaller the number of channels. But if you reduce the permitted power for each channel, thus reducing S_d, then all users within the cell will suffer. This might prevent an access point from communicating on a particular channel with a station because it would be out of range, even though both of them were well inside the cell boundary.

One approach to minimising this problem might be to broadcast a number of channel adverts, at different strengths. Consider a channel C. The beacon would first advertise the availability of this channel at low power. This would reach those users within region R_1 in Figure 4. The advert would allow the channel to be used at power P_1, with maximal range D_1. The same channel would then be advertised at maximum power, but with permission to use the channel at a lower power setting

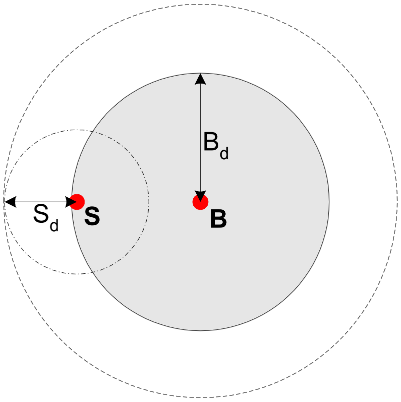

Fig. 3. Fixed power cell size

D_2. A device receiving multiple transmissions from a beacon would be allowed to use the highest power setting. This technique would reduce the size of S_d, whilst allowing devices closer to the beacon to still run at full power, as illustrated in the figure.

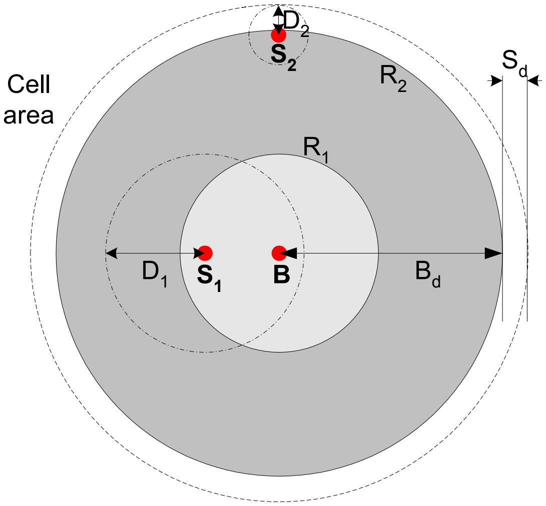

Fig. 4. Variable power cell size

Client radios, e.g. 802.11 access points and stations, would use the transmissions from the nearest or strongest beacon to determine which bands and channels are currently available for use. When out of range of an authenticated beacon they would just use a set of default channels, e.g. the current set of 802.11b channels. The information transmitted by the beacon would be signed to avoid the possibility of a bogus beacon being introduced into the system. Ideally you would like the communication mechanism and protocol between client and beacon to be such that a relay device couldn't be constructed to extend the range of the beacon without being detected. Perhaps the simplest strategy would be to have a challenge/response exchange that has to be completed within a limited time window. This wouldn't completely prevent a relay being constructed, but could perhaps prevent this being accomplished without specialised hardware. The digital certificate of the beacon gives us a guarantee of the beacon's location. If we can be reasonably certain that the client machine is within a certain distance from the beacon then this location service could have other potential uses, e.g. increased access control on connections to a corporate network.

The process used for establishing a communication between an access point and a user station would need to change to make use of these facilities. We assume that access points would still broadcast beacon frames on the standard 802.11 channels, and this mechanism would be used to establish initial contact between an access point and a station. At this point both devices will have received a set of channels and power levels they can use from the frequency beacon advert(s). This information may not be identical in both devices for a variety of reasons. If the devices are at different distances from a beacon then the device furthest away may have not received some of the adverts (Figure 5a). The two devices may even be serviced by different beacons, where cells overlap for example (Figure 5b). So a negotiation phase needs to take place to determine an appropriate channel and associated power level. Note that it is not sufficient to just pick a channel they have in common. A station near the edge of the cell may only be able to use a channel at low power, whereas the access point may be allowed to use the channel at a higher power level if it is nearer the beacon. This might result in the access point being able to contact the station on this channel, but the reverse direction might fail. Different channels may have different power constraints, depending on what else uses this band, and so in such a case another channel would be chosen, for one or both directions.

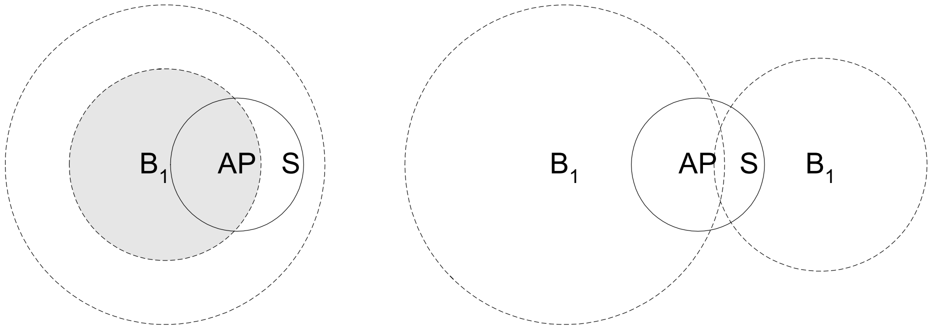

Fig. 5. a) Different ranges b) Different beacons

The situation where WiFi devices are used outside the range of a beacon is also complicated. It is tempting to just use the standard channels in such cases, but as we have already seen these channels depend on the country you are in. One approach might be to use a hierarchy of beacons. One beacon would cover the whole country, broadcasting the standard channels for that country. Localised beacons would then broadcast the more location-specific constraints for those areas with enough users to justify the complexity. The data rates required for the beacon transmissions would be fairly minimal, and the access points and stations would only need to be able to receive on these bands, not transmit. It should therefore be relatively straightforward to identify a small set of bands to use for such transmissions that could cover the entire world.

Stations sometimes incorporate mechanisms to vary their transmitter power depending on the estimated range from the receiver, or the observed error rates. Such mechanisms are required for increased spectrum efficiency and longer battery life. These mechanisms would be unaffected by this proposal other than the requirement to observe the power limits advertised by the beacons. In some situations the power control mechanism might suggest using a power level that is prohibited in the current geographic location. In these cases the station would need to move to a different channel, with a more liberal power policy.

Minimising battery usage is a crucial aspect in the design of many devices. Any scheme that required the receiver to be always on, or the transmission of large numbers of additional packets, would be at a severe disadvantage. The current proposal is fairly well-behaved in this respect. No device other than the frequency servers has to transmit back to the beacons. Connection establishment between an access point and a station may require more exchanges to negotiate a channel to use. Roaming will be slightly more complex as well, as you may need to renegotiate a channel, even to the same access point, when you move out of range of a low-power channel. The proposal doesn't stop a device going to sleep. The only requirement would be that when it woke up again it would need to wait for a signal from the beacon confirming the right to use the channel before transmitting on the channel again.

The proposal does not address the issue of automating cell-site planning. The beacons are not aware of any devices in their coverage area. They are merely broadcasting the right to use particular channels within their area. Any device wishing to use a "non-standard" channel would have to be "beacon-aware", not just access points. This proposal doesn't automate channel selection in any way. In fact it makes it harder, as there is now more choice, and this choice could vary over time. But even without the frequency server approach it seems clear that we will eventually reach the stage where an access point may have many channels and bands to choose from, and be able to use multiple channels simultaneously. In such a setting it will become unacceptable to rely on a user configuring these choices manually, particularly when different access points are under the control of different people. So peer-to-peer protocols between access points will need to be deployed to negotiate channel usage policies for the access points that minimise interference.

Mobility is always an issue in wireless networks. It is clear that the frequency server/beacon approach would not work, at least as currently proposed, for a "high-velocity" mobile device such as a car. It should work for a user roaming around a site though. If the beacons transmitted the channels once a second, or that order of magnitude, then a user isn't going to get far between transmissions, and so the "guard band" will be quite small. As with the initial session establishment, the handover between access points would have to be a bit more complex to negotiate an appropriate channel to use. But some of this complexity will also occur if access points start to use multiple channels anyway, and so the additional overhead is perhaps quite minor.

It would be tempting to just advertise wide chunks of spectrum to use. But this would raise a number of complex usage issues and so it is safer to assume the beacons will advertise channels rather than bands. Signal bandwidth and ACP considerations would be taken care of when the additional channels were "designed". One possibility enabled by this scheme is that channels could be allowed to overlap. For example, in the 5 GHz band you could imagine allocating a bunch of 802.11b channels that occupied the same spectrum as the current 802.11a channels. You clearly couldn't use both at once, but if on a particular site you only wanted to use 802.11b, not 802.11a, then the beacon would just permit the additional channels, blocking the use of the 802.11a ones.

4. CONCLUSIONS

The frequency server approach obviously has a lot more infrastructure than an autonomous system that just listens for unused frequency bands. However, it also has a number of advantages. Depending on the range of the beacons, and the sophistication of the servers, the beacon approach could allow a very fine-grained control over frequency usage, both in the geographic and temporal dimensions. You could use frequencies allocated to the fire brigade within a building, and then disable this usage when a fire alarm went off, for example.

Careful packaging, to prevent easy disassembly of the components, plus digital signing could give us fairly good guarantees that restricted areas would be protected. The aim would be to make it take as much effort to work around the mechanism as it would to build an illegal radio. The beacon/server infrastructure could also be exploited for other uses, spreading the costs over a wider range of uses. The ability to determine an approximate position for a machine, together with cryptographic confidence that this information is reliable, could be very useful in many areas. Although our discussion has focused on supporting WiFi, the same infrastructure could be used for other services. For instance the UK has a dedicated band allocation of 183.5MHz to 184.5MHz for Remote Meter Reading applications. Other countries may allocate other bands for such devices. It could eventually become cheaper to make a device "universal", using the frequency server transmissions to determine the band to use for such a device, rather than making country-specific versions, where there is always the risk of such devices being used in countries for which they were not intended.

Pushing the boundaries further, you could even imagine the beacons being able to download new firmware to wireless stations, e.g. if a particular country used a modulation scheme that the device was unfamiliar with, assuming we were using SiRiFIC or some other kind of "soft" radio. Commissioning a beacon need not be a time-consuming process. For example, there might be a standard profile for internet cafes in a high street setting. In many cases this could avoid an expensive site survey and may be sufficient if all that was required was access to a small set of additional channels at low power.

Clearly the legislating authorities would need to have a lot of confidence in the design and robustness of the frequency servers and beacons before they would authorise the use of such a technique. As mentioned earlier, a combination of cryptography and design accomplishes much of this, but trusted manufacturers would also have a role to play in building confidence. WiFi manufacturers would then have to demonstate their devices respected and authenticated the beacon transmissions before being licensed to use the wider range of bands.

REFERENCES

Droms, R. 1997. Dynamic host configuration protocol. Tech. Rep. RFC 2131, Internet Engineering Task Force. March.

Kwerel, E. and Williams, J. 2002. A proposal for a rapid transition to market allocation of spectrum. Tech. Rep. OPP Working Paper 38, Federal Communications Commission. November.

Mockapetris, P. 1987a. Domain names – concepts and facilities. Tech. Rep. RFC 1034, Internet Engineering Task Force. November.

Mockapetris, P. 1987b. Domain names – implementation and specification. Tech. Rep. RFC 1035, Internet Engineering Task Force. November.

Motorola. 2002. The exploitation of “spectrum holes” to enhance spectrum efficiency. Tech. rep., Motorola. October.

SiRiFIC. 2003. Wireless single chip multimode transceiver technology. http://www.sirific.com/.