Managed Unlicensed Spectrum

William Webb, June 2017

Introduction

Unlicensed spectrum is only very loosely managed. Rules of access such as maximum power levels and duty cycles aim to regulate behaviour to a limited degree but by necessity are set in place before a band is put into use. Database approaches such as those proposed for TV white space can be more proactive by dynamically changing whether access is allowed but do not seek, for example, to minimise interference between unlicensed users.

For many, the lack of management is one of the key attractions. With little control over what can be deployed, innovation can occur and spectrum access is enabled for all. Most would judge unlicensed spectrum to have been a great success with widespread and growing usage of Wi-Fi, Bluetooth, myriad connected home devices and much more. Hence, the idea of imposing some kind of management control might be seen as contrary to one of the founding principles on which unlicensed access has been built.

In this paper, we look at whether there may be some cases where management is appropriate and some mechanisms that can be deployed to achieve it without compromising the attributes that have made unlicensed spectrum so attractive.

Why management might be appropriate in some cases

Broadly, management of spectrum is only appropriate where interference is a problem. If there is no interference then there is no need for management. If we look at some unlicensed usage we might categorise interference problems broadly as follows:

Interference not a problem	**Interference can be problematic**
Cordless telephones (own spectrum)	Wi-Fi in dense areas
Wi-Fi in less dense areas	Wide-area IoT solutions
Bluetooth	Drones
Home connectivity	RFID
Car door openers	
Wireless keyboards, mice, etc	

Simplistically, interference tends not to be an issue where one of the following holds true:

- The usage has its own unlicensed spectrum (eg cordless phones in Europe).
- The density of usage is low, such as outside urban areas.
- The solution is highly tolerant of interference, such as Bluetooth where frequency hopping provides excellent interference rejection.
- The usage is very short range such as wireless keyboards.

Interference tends to become problematic where:

- The density of usage is high, such as Wi-Fi in conference venues.
- The range required is long, such as wide-area IoT solutions with a range of 5km or more.
- The power levels are very low, such as RFIDs, where any interference is problematic.

This paper concentrates on two case studies to demonstrate the benefits of a solution and show what might be implemented before drawing some general conclusions.

1. The use of Wi-Fi where uncoordinated frequency planning between Wi-Fi nodes can halve the spectrum efficiency, and where new technologies such as unlicensed-LTE might cause future problems.
2. Wide area IoT connectivity where many uses are unlikely because of concerns over being able to guarantee a quality of service level in unlicensed spectrum.

Wi-Fi

For most people, their experience of broadband connectivity is via a Wi-Fi connection. In the home almost all computing devices, tablets and smart-phones are Wi-Fi connected rather than being plugged into a wire connected to the home gateway. Approximately 80% of all our mobile data consumption is via Wi-Fi and over 80% of all our computing devices have Wi-Fi connectivity but no cellular connectivity.

In some cases, Wi-Fi can form the bottleneck in the delivery of higher speed content. If the connection to the home is 100Mbits/s but the Wi-Fi connections only support 10Mbits/s due to congestion, interference or weak signals then the higher speed of access to the home is of limited value. In congested areas such as apartment blocks, shopping malls, train stations and conference centres, the level of interference between Wi-Fi nodes can be such that data rates fall below 1Mbits/s. If Wi-Fi congestion were to rise then many attempts to deliver faster broadband would not succeed. With ever more Wi-Fi nodes deployed then a rise in congestion is very likely.

Reduced data rates in Wi-Fi can be caused by three effects:

- Low signal level caused by being far from the router.
- Congestion caused by many devices accessing the same router.
- Interference between routers.

The solution to the first two problems is to install more routers to provide greater coverage and capacity, but as more routers are installed interference rises. Hence, reducing interference between routers is critical to enabling improvements in Wi-Fi data rates.

In cellular systems, central planning of the radio frequency assigned to cells ensures neighbouring cells do not interfere and that devices are attached to the optimal cell. In most Wi-Fi systems this decision making is decentralised with each router selecting what it believes is the best frequency and each device attaching to its preferred router: often according to whether it has attached to that particular one in the past. This decentralised approach works well when router density is low and each router is able to find a free channel. However, it works poorly when there are insufficient channels and in some cases algorithms may actual make matters worse as routers “fight” against each other for optimal channels.

The solution is some degree of centralisation of the management of router frequencies and device selection of routers in those areas where there is a high density of routers. This centralisation could be local, covering only a single building or dense area, or it could be regional or national. However, this centralisation is complicated by the fact that routers are owned by many different parties. For example, in a block of flats, each flat owner will likely own a router. In a train station, each of the shops and restaurants may own their own routers. In a city centre there may be a mix of routers owned by mobile operators, by third-party Wi-Fi providers, by city councils and by office tenants. None of these may wish to take the lead in providing centralised planning. In this situation a third-party or agreed shared mechanism is needed that can control the access points, the device attachment and deliver an optimised service for the benefit of all. Achieving this requires:

- Standardisation such that each router knows how to contact the centralised manager, is able to provide information in a standard format on frequency usage and the manager knows how to direct the router to act.
- The emergence of a trusted third party or trusted shared solution to provide this management function.
- Appropriate central planning solutions that deliver optimal results.
- Wi-Fi frequencies as free as possible from interference from non-Wi-Fi devices since interference that could not be controlled would complicate the central planning process.
- The participation of all, or almost all, routers in an area in the scheme – if only a small subset participate there will be limited benefits[1].
- An appropriate business model to cover any costs that arise, complicated by the fact that few currently pay for delivering Wi-Fi and that there are few immediate advantages for the first to join the scheme.

Elements of many of these already exist. The standard TR-069 enables Wi-Fi nodes to send information to a remote management unit and to receive instructions in return. Companies such as Vodafone already manage the routers that connect to their broadband lines using this standard, planning frequencies throughout homes. Some airports manage Wi-Fi frequencies, requiring franchise owners on the airports to link their Wi-Fi routers to a central unit. Software solutions that can be embedded within routers and central planning algorithms have been suggested and trialled by commercial and academic entities. But operation at scale across routers owned by different parties is rare and would require significant leadership to achieve.

For example, EcoWi-Fi – a commercial venture – claims to be able to improve Wi-Fi data rates by up to 130%. Aoifes, the company developing EcoWi-Fi, is hoping to partner with ISPs and router manufacturers, targeting home deployments of Wi-Fi routers. Their solution comprises additional software embedded in routers and a cloud-based central management system. Another approach is termed Empaticradio and has been developed by academics in Norway and aims for peer-to-peer communication between routers with limited need for centralised management. Instead routers aim to decide between themselves as to the best approach. Alternatively, organisations such as Assia have proprietary solutions that aim to optimise the performance of individual routers in the presence of interference, claiming over 100% performance improvement[2].

Concerning the emergence of either a trusted third party to manage the central database or a peer-to-peer solution there appears to be two significantly different approaches:

- A centralised database owned by a legal entity (eg a company, Government) that registers routers and manages their key parameters.
- A peer-to-peer solution where routers self-discover other nearby routers and using cloud-hosted software collectively determine their optimal frequency allocation.

At this stage there is no clear preference between these – and indeed both could potentially co-exist as long as there was some cooperation between them. If there is central management then in the

[1] It is not clear as to what percentage of routers would need to participate to realise the majority of the benefits. There is anecdotal evidence that it is not necessary to achieve 100% but further study is required to understand the minimum percentage that would make the undertaking worthwhile.

[2] Simulations performed by the author suggest similar levels of gains in eg dense residential areas.

first instance confidence will be improved if a third party is sanctioned by Government, a regulator or some similar entity. This could be as part of a procurement process to provide management to Government-owned routers.

Where there is a cost then options might include:

- Owners of Wi-Fi routers paying a monthly fee for the central planning approach.
- Government paying the central planner for the benefits it provides, perhaps from revenues arising from spectrum auctions, fee payments or similar.
- A small fee being attached to fixed line usage to cover central Wi-Fi planning since most fixed lines will have an attached Wi-Fi router.
- Router manufacturers funding the third party through a small increase in the cost of routers.

The business model would need to ensure that it sufficiently incentivised those organisations that already manage groups of routers to partake in a wider management process. The peer-to-peer approach may have fewer costs making the business model less critical, but key players will still need some form of incentivisation.

Implementing a solution of this sort is complicated by the fact that there are already many deployed routers and that it might not be possible to locate and update these remotely. Hence, it may require gradual churn and replacement of devices. This means there is very little benefit from early adopters so if they have to pay any fee they may decide to wait until the majority of devices have enrolled into the solution. If all make this decision, then the solution is never successful.

Further study is required to determine the best way to resolve this start-up problem. Some areas that might merit investigation include:

- Opening a new band, or partial band (eg at 5GHz) and requiring all routers using this band to have central management.
- “Seed-funding” the process such that Government or others provide the funding needed to enable the introduction of the solution to the point that the benefits were clear. (At this point the initial funding could potentially be repaid.)
- Local initiatives such that, eg, all the residents in a particular area were encouraged to enrol in a relatively short timeframe.

Internet of Things

IoT connectivity broadly divides into short-range in the home, office or factory; and long-range for devices that need connectivity outdoors. The latter include systems like smart metering, connected trash-cans, asset tracking, smart parking sensors, agricultural sensors and much more.

Current solutions to long-range or wide-area connectivity fall into two categories often referred to as licensed and unlicensed. Licensed solutions are those deployed by the mobile operators in their licensed spectrum. These have been developed within 3GPP and comprise LTE-M and NB-IoT. Simplistically, LTE-M is more appropriate for devices with substantial batteries or mains power, while NB-IoT is better for the low-power and lower-functionality devices. Some mobile operators, such as Vodafone, are planning aggressive roll-out of licensed solutions across their network in 2017 and 2018.

Unlicensed solutions have predominantly been developed as proprietary technologies by companies such as Sigfox or the LoRa technology developed by Semtech, or standardised technologies such as that from the Weightless SIG, and are deployed by a wide range of different entities in unlicensed spectrum. Some deployments are campus-wide self-provision, whereas others are national networks deployed by operators such as Sigfox and their affiliates. Whether both licensed and unlicensed solutions will continue to exist and the balance of traffic between them is unclear. Some believe that the mobile operators have such a compelling advantage in terms of coverage and branding that they will dominate the market. Others think that the relatively high cost of cellular-compatible modules and cost-structure of the operators will result in licensed solutions only being used for the most valuable of connections. Some look to the world of personal connectivity where a mix of licensed cellular solutions and unlicensed Wi-Fi solutions are in use.

Predicting the outcome appears near-impossible at this point. However, there is a risk of a default to a sub-optimal outcome if both licensed and unlicensed solutions are not available for operators and users to select from.

One of the key problems for unlicensed deployments is a lack of suitable spectrum. Most unlicensed spectrum is configured for short-range devices which communicate infrequently. Therefore, it has limits on transmitted power and on duty-cycle, often as low as 1%. A wide-area IoT solution typically comprises a base station which needs to transmit for somewhere between 10% and 50% of the time and could benefit from using relatively high power levels, although terminal devices generally will only need 1% transmission time and be low-power. However, relaxing the duty cycle and transmit power could increase interference[3].

A second problem is that the compromises made in system design to enable low-power devices to communicate over many kilometres, can make the systems vulnerable to interference. For example, recent reports suggest that the LoRa solution can experience a rapid reduction in performance as interference from other technologies increases. Some solutions may fail completely if that have no way of feeding back to the device the need for a change in parameters due to interference[4].

A solution to these issues would be a different set of spectrum-access rules for wide-area IoT solutions. Ideally, these would be accompanied by spectrum dedicated to unlicensed IoT operation. For example, higher powers and unlimited duty cycles could be allowed if operators agreed to central coordination which delivered a fair division of the spectrum across networks in a manner that optimised their performance. This might segregate technologies known to have poor co-existence while frequency planning technologies that can co-exist.

Not only would this enable the deployment of networks that might otherwise not be possible, it would go some way towards reassuring potential users of the network that significant interference issues would not occur in the future, especially if dedicated bands were provided.

The management probably needs to be centralised – unlike Wi-Fi where peer-to-peer management might be viable. It could be Government owned, commercial or run by a not-for-profit standards body or similar. Possible models include dynamic spectrum sharing databases. Because this is an emerging area and because the number of players is relatively small then the start-up problem is minimal and the business case unlikely to be an issue.

[3] There are compromise positions – for example Ofcom allowed a 10% duty cycle for base stations where the location was registered with the regulator.

[4] For example, the Sigfox solution has no significant downlink and hence no way of changing device behaviour. Devices will continue to transmit regardless of whether their communications have any chance of being received.

Implications and conclusions

This paper has suggested that while managed unlicensed spectrum is not needed nor appropriate in many cases, there are important situations where it could make a material difference. Indeed, given the critical role of Wi-Fi and the future potential of IoT to transform society and address some of our key challenges, even small differences in spectrum efficiency or increases in potential use cases could deliver huge societal benefits.

Where it is needed, there are different issues and ways to achieve solutions. In Wi-Fi, peer-to-peer self-organisation might be viable, in IoT changes in spectrum access parameters when utilising a centralised database might be preferred. Some of the key attributes of the two solutions are summarised in the table below.

Parameter	**Wi-Fi**	**IoT**
Voluntary or obligatory?	Voluntary, although the start-up/free-rider problem needs addressing	Obligatory if enhanced spectrum access to be allowed
National or international?	Could work on a national basis but given the international nature of equipment supply, best addressed internationally	National, set by each regulator, although economies of scale helped if multiple regulators adopt same approach
Getting started	Could happen without Government or regulatory involvement, eg by Wi-Fi Alliance	Needs regulatory lead
Timescales	Could be prolonged	Needs to ideally happen in the short-term

There are strong grounds for trialling managed unlicensed solutions in both these cases in order to assess the role management might play in the future.

Biography

William is the CTO at OpenSignal, the world-leader in crowd-sourcing of mobile connectivity information. He is also CEO of the Weightless SIG, the standards body developing a new global M2M technology. He was President of the IET – Europe's largest Professional Engineering body during 14/15.

He was one of the founding directors of Neul, a company developing machine-to-machine technologies and networks, which was formed at the start of 2011 and subsequently sold to Huawei in 2014 for $25m. Prior to this William was a Director at Ofcom where he managed a team providing technical advice and performing research across all areas of Ofcom's regulatory remit. He also led some of the major reviews conducted by Ofcom including the Spectrum Framework Review, the development of Spectrum Usage Rights and most recently cognitive or white space policy. Previously, William worked for a range of communications consultancies in the UK in the fields of hardware design, computer simulation, propagation modelling, spectrum management and strategy development. William also spent three years providing strategic management across Motorola's entire communications portfolio, based in Chicago.

William has published 15 books, over 100 papers, and 18 patents. He is a Visiting Professor at Surrey and Southampton Universities, an Adjunct Professor at Trinity College Dublin, a Board member of Cambridge Wireless, a member of the Science Advisory Council at DCMS, other oversight Boards and a Fellow of the Royal Academy of Engineering, the IEEE and the IET. In 2015 he was awarded the Honorary Degree of Doctor of Science by Southampton University in recognition of his work on wireless technologies and Honorary Doctor of Technology by Anglia Ruskin University in honour of his contribution to the engineering profession. His biography is included in multiple "Who's Who" publications around the world. William has a first class honours degree in electronics, a PhD and an MBA.