Weightless: The technology to finally realise the M2M vision

Professor William Webb *FREng, FIEEE, FIET*

Abstract

The value in machines having wireless communications has long been understood and a large market predicted for many years. That this has not transpired has been because of the difficulty of meeting all the requirements for machine communications within the constraints of the available radio spectrum. These constraints changed significantly with the advent of white space which provides near-perfect spectrum with free access. However, the combination of the unique and unusual nature of that access and the very different characteristics of machine traffic compared to human traffic means that using any existing standard is far from optimal. This paper discusses the design of a new standard called Weightless which is designed and optimised for this application.

Keywords: White space, radio spectrum, M2M, machine communications, Weightless.

1 Machine communications is lacking spectrum to enable a ubiquitous network

Machine communications, often termed M2M, has long been forecast to be a sector with massive growth. Over the last few decades many have noted that the installation of a wireless connection into myriad devices would bring a range of benefits. An enormous range of examples have been suggested, from cars to sensors to traffic lights to healthcare applications and much more. More recent forecasts of 50 billion connected devices by 2020 do not sound so incredible when it is realised that this is only ten devices per person – many people already have more than ten wirelessly enabled devices in their home.

However, the market for machine communications to-date has been weak. There are some cars with embedded cellular modems and some relatively high-value items such as vending machines are equipped with cellular packet-data modems. But the market today is only a tiny fraction of the size it has long been predicted to grow to. This is predominantly due to the lack of a ubiquitous wireless standard that meets the needs of the vast majority of the machine market. These needs include:

- Low cost, both of the hardware and the service. Many machines are individually of relatively low value – imagine for example a temperature sensor. Chipset costs need to be in the region $1-$2 and annual service charges less than $10 to make it worth embedding wireless technology in such devices.
- Excellent coverage. To make applications such as smart metering viable there needs to be coverage of near 100% of all meters. With many meters deep within the home or even in basements this implies vastly better coverage than achieved with today’s cellular networks.
- Ultra low-power operations. Many machines are not connected to the mains and so have to operate on batteries. Having to change the battery is at best an annoyance and at worst a significant expense. Battery life of ten years or more is essential.
- Secure and guaranteed message delivery. While machines rarely need ultra-rapid transmission, they do need to be certain that messages have been received and that security of the system has not been compromised in any way.

There is no current wireless system that comes close to meeting all of these requirements.

Cellular technologies do provide sufficiently good coverage for some applications but the hardware costs can be $20 or more depending on the generation of cellular used and the subscription costs are often closer to $10 per month than $10 per year. Battery life cannot be extended much beyond a month. Cellular networks are often ill-suited to the short message sizes in machine communications resulting in massive overheads associated with signalling in order to move terminals from passive to active states, report on status and more. So while cellular can capture a small percentage of the market which can tolerate the high costs and where devices have external power, it will not be able to meet the requirements of the 50 billion device market. Indeed, if it could, it would have done so already and there would be no further debate about the need for new standards.

There are many short-range technologies that come closer to the price points. These include Wi-Fi, Bluetooth, Zigbee and others. However, being short range these cannot provide the coverage needed for applications such as automotive, sensors, asset tracking, healthcare and many more. Instead, they are restricted to machines connected within the home or office environments. Even in these environments there may be good reasons why a wide-area solution is preferred. For example, an electricity supply company is unlikely to accept that their meter is only connected via, eg Zigbee, into a home network, which in turn connects to the home broadband. Were the home owner to turn this network off, fail to renew their broadband subscription or even just change the password on their home router, then connectivity could be lost. Restoring it might require a visit from a technician with associated cost. Maintaining security across such a network might also be very difficult.

Finally, it is critical that the technology is an open global standard, rather than a proprietary technology. With a wide range of applications there will need to be a vibrant eco-system delivering chips, terminals, base stations, applications and more. The manufacturer of a device such as a temperature sensor will need to be able to procure chips from multiple sources and to be sure that any of them will interoperate with any wireless network across the globe.

Without a wide-area machine communications network that meets all of the sector requirements it is unsurprising that forecasts for connected machines have consistently been optimistic.

While the needs of the machine sector have long been understood, the key problem to date has been a lack of insight as to how they could be met. Ubiquitous coverage requires the deployment of a nationwide network, and the conventional wisdom has been that such networks are extremely expensive. For example, a UK-wide cellular network can readily cost $2 billion with costs of spectrum adding another $1-2 billion. With the machine market unproven, such investments were not justifiable and would result in an overall network cost that would not allow the sub $10/year subscription fees needed to meet requirements.

The key to unlocking this problem is free, plentiful, globally harmonised low-frequency spectrum. It needs to be free, or at least very low cost, to keep the investment cost low. It needs to be plentiful to provide the capacity to service billions of devices. It needs to be globally harmonised in order to allow devices to roam across countries and to enable the economies of scale needed to deliver <$2 chipsets. It needs to be low-frequency to enable good range from each base station and therefore a relatively small number of base stations to provide ubiquitous coverage. Unfortunately, low frequency spectrum is in very high demand and so rarely become available in sizable quantity, is almost never globally harmonised and where even a few of these attributes hold true is extremely expensive.

The lack of spectrum that meets all these requirements has meant that up until now the only option for wide-area machine communications has been to make use of existing networks, predominantly cellular.

2 White space as a key enabler

In the last year, a new option has emerged for spectrum access [Gurney 2008, Nekovee 2010]. This is the use of the “white space” spectrum – the unused portions of the spectrum band in and around TV transmissions. White space meets all of the requirements for M2M communications. It is unlicensed and so access to it is free. It is plentiful with estimates of around 150MHz of spectrum available in most locations – more than the entire 3G cellular frequency band. It has the potential to be globally harmonised since the same band is used for TV transmissions around the world. Finally, it is in the perfect low frequency band which enables excellent propagation without needing inconveniently large antenna in the devices. This is a “game changer”. Access to white space provides the key input needed to make the deployment of a wide-area machine network economically feasible.

However, white space is not without its issues. These are broadly regulation and interference.

Regulation for white space is still developing in many countries [Ofcom,FCC,DVB]. However, it is clear that white space access will require devices that have the following characteristics:

- Relatively low output power. The FCC has specified 4W EIRP for base stations and 100mW EIRP for terminals. These are an order of magnitude lower than cellular technologies.
- Stringent adjacent channel emissions. White space devices must not interfere with existing users of the spectrum, predominantly TVs. Hence, the energy that they transmit must remain almost entirely within the channels they are allowed to use. The FCC has specified that adjacent channel emission need to be 55dB lower than in-band emission, a specification much tighter than most of today’s wireless technologies.
- The need to frequently consult a database to gain channel allocation. Devices may need to rapidly vacate a channel if it is needed by a licensed user. They must consult a database to be informed as to the channels they can use and must quickly move off these channels as required.

Interference can be problematic in white space. Many channels have residual signals from TV transmissions. These can either be in-band emissions from distant, powerful TV masts that are too weak for useful TV reception but still significantly above the noise floor. Alternatively, they can be adjacent channel emissions from nearby TV transmitters some of which are transmitting in excess of 100kW. In addition, since the band is unlicensed, other users might deploy equipment and transmit on the same channels as the machine network, causing local interference problems.

These are not insurmountable issues. But no current technology has been designed to operate in such an environment and so would be sub-optimal at best. For example, we have shown that in the UK an optimised technology could access around 90MHz of white space after all the interference issues are taken into account, whereas an existing technology such as Wi-Fi or WiMax could only access around 20MHz.

Hence white space spectrum provides the key to unlock the machine network problem. But it comes at the cost of needing to design a new standard.

3 Design rules for an M2M solution

While the use of white space provides the need for a new standard, there are many benefits from designing a standard specifically for machine communications. Machines are very different from people. Typically, their requirements vary in the following manner:

- Much shorter message size than most human communications (with the exception of SMS text messages). Most machines only send a few bytes of information whereas as person may download Mbytes of information.
- More tolerant of delay. Most machine communications is relatively unaffected by a few seconds of delay whereas people quickly find this frustrating.
- Generally predictable communication patterns. Machines often send data at regular intervals and so can be “pooled” on these occasions. People’s communication needs are typically unpredictable and so contended access for resources is needed.

Taking advantages of these differences allows the design of a system that is much more efficient, providing greater capacity than would otherwise be the case and hence having low cost. The predictability of most communications allows a very high level of scheduled communications as opposed to unscheduled, or contended, communications. The difference is akin to pre-booking passengers on flights so that each flight is full, but not over-crowded, rather than just letting passengers turn up, as with most trains, and suffering the crowding problems that occur. By telling terminals when their next communications is scheduled, future frames of information can be packed very efficiently and terminals can be sent to sleep for extended periods extending battery life.

Scheduling brings many other advantages. The first is efficiency. Contended access schemes can only operate up to about 35% channel usage – above this level the probability of access messages clashing becomes so high that very little information gets through. By comparison, scheduled access can achieve close to 100% efficiency. Scheduling can be enhanced by complex algorithms in the network that prevent terminals close together in neighbouring cells transmitting simultaneously, that ensure terminals suffering local interference are scheduled on frequency transmissions where interference is minimised, and much more.

Another design rule for M2M is that coverage is typically more important than data rate. For example, it is more critical that all smart meters can be read than what the data rate of transmission is – as long as it is sufficient to transfer data regularly. In fact, most machine communications can be measured in bits/s rather than kbits/s or Mbits/s. As an example, a smart meter will typically send around 20-40 bytes of information perhaps once every 30 minutes. This equates to an average of 240bits per 30 minutes or 8 bits/minute. There are applications that will require higher data rates, but speed is rarely critical. Hence, a good M2M system design will trade off data rate against range. This can be achieved by spreading the data to be transmitted. Spreading involves multiplying the data by a pre-defined codeword such that one bit of transmitted data becomes multiple bits of codeword. The receiver can then use correlation to recover the codeword at lower signal levels than would otherwise be possible. Codewords are selected to have particular correlation properties and typically have length 2^n (eg 16, 32, 64). So, for example, multiplying the transmitted data by a codeword of 64 results in an improvement in link budget of some 18dB but reduces the data rate by a factor of 64. Most buildings have a penetration loss for signals entering them of around 15dB so spreading by this factor would provide indoor coverage to machines where only outdoor coverage previously existed. Some M2M solutions have spreading factors extending as far as 8192, providing great range, but very low capacity. Large spreading factors do add complexity to the system design since they extend the time duration of important system control messages that all devices must hear, which in turn requires long frame durations. These design decisions make M2M networks radically different in many respects from cellular solutions.

Another design rule, at least at this embryonic stage of the market, is flexibility. It is far from clear what M2M applications will emerge. Even the balance between uplink and downlink is unclear – for example smart meters will likely generate predominantly uplink traffic while software updates, perhaps for car engine management systems, will be large downlink messages. This suggests that systems should be time division duplex (TDD) in order that the balance between downlink and uplink can be changed dynamically.

M2M systems should make the terminal as simple as possible, keeping complexity within the network. This is contrary to the trend in cellular communications where handsets have been becoming ever more powerful and complex. There are two key reasons to keep M2M terminals simple. The first is to keep the cost as low as possible – as mentioned earlier many applications require chips with costs of the order $1-$2. The second is to minimise power consumption for terminals that are expected to run off batteries for 10+ years. This means that, for example, complex multi-antenna solutions should be avoided and that terminals should not be expected to make complex calculations to decode their messages. Even an apparently simple decision, such as requiring a terminal to respond on the uplink of a frame where it receives a message on the downlink could require it to process the downlink message much more rapidly, needed a more powerful processor. Careful design throughout is needed to achieve minimal terminal complexity.

Finally, there is likely to be an imbalance within an M2M network where the base station has much more power and processing at its disposal and so can have a greater range than the terminals. This is of no value since the terminals need to be able to signal back and so the link budget must be balanced. With base stations transmitting often around 4W (36dBm) but battery powered terminals restricted to 40mW (16dBm) there is a 20dB difference. This can be balanced by a combination of using narrower bandwidths in the uplink and using greater spreading factors.

Designing M2M solutions does not require any technological break-through. But it does require great care in understanding the implications of each decision and it needs a system design that is radically different from a cellular network, with design decisions often appearing contrary to the conventional wisdom of the day.

4 Design rules for white space

White space is unique spectrum. It is the first band where unlicensed users are allowed to mix with licensed users as long as they do not cause any interference to those users [Wang 2011,Fitch 2011]. This brings the benefit of free access to highly valuable spectrum but also a need to operate in an uncertain environment. Any system operating in white space should adhere to the following design rules.

Firstly it needs very low levels of out-of-band emissions. This minimises interference caused to licensed users and so maximises spectrum availability. Achieving such low emission levels means that modulation schemes such as OFDM should be avoided as these tend to have relatively large adjacent channel emissions.

Next it needs to avoid interference caused by other unlicensed users which can be random and sporadic. Classic techniques for doing this include frequency hopping to rapidly move off compromised channels. However, hopping in a network requires central planning to avoid neighbouring cells using the same frequencies. Optimal planning where different frequencies may be available in different cells and the sequence may need to dynamically adapt to interference is complex and requires new algorithms.

Where interference cannot be avoided the system needs to be able to continue to operate. Spreading, as discussed above, can also be useful to work in channels with interference, again trading off range (or tolerance to interference) against data rate. Base stations can often experience significant interference from distant TV transmissions and require mechanisms such as interference cancellation to reduce its impact.

Finally, where there are few white space channels available, it can often be possible to increase availability by transmitting with lower power and hence causing less interference. Power control is therefore critical, again coupled with spreading where needed to regain the range lost from the lower power.

5 Weightless – the standard designed for M2M in white space

Designing *the* standard for M2M in white space requires many trade-offs and iterations. A key starting point is the conflict between excellent coverage requirements and yet low-power constraints both due to white space regulation and the need for long battery life in terminals. The only way to achieve long range with low power is to spread the transmitted signal. Hence, variable spreading factors from 1 (no spreading) to 1024-fold are a core part of the Weightless specification. Spreading is essentially a mechanism to trade range against throughput so using high spreading factors can achieve significant range extension but at the cost of lower data rates. Happily, there is sufficient bandwidth in the white space frequencies, and M2M data rates are sufficiently low that more than adequate capacity and throughput can still be achieved even with high levels of spreading.

Use of the white space spectrum does not provide guaranteed uplink and downlink pairing, making TDD operation essential. This in turn leads to a frame-structure with a downlink part then an uplink part which repeats periodically. The maximum spreading factor informs what this repetition should be since the header information at the start of the frame needs to be spread by the maximum factor in order that all terminals in the cell can decode it. If this header takes up more than around 10% of the frame length then the system starts to become inefficient as signalling becomes a significant percentage of the total traffic. Simple calculations show frame lengths of around 2s are optimal. This would be overly long for person-to-person communications, with such a delay being highly annoying, but is not an issue for M2M communications (machines do not generally get annoyed!).

The need for stringent adjacent channel emission levels suggest the use of single-carrier modulation (SCM) rather than OFDM as the latter is more difficult to filter tightly without distorting the transmitted signals. OFDM also has a high peak to average power ratio which does not fit well with very low powered devices. Because the terminals are often very low power (eg 40mW) compared to base stations (which can be up to 4W) the link budget needs to be balanced. This is achieved with a narrower band uplink such that the noise floor is lower. Using around 24 uplink channels for each downlink has the effect of balancing the link budget.

Operation in white space requires good interference tolerance. This is achieved primarily using frequency hopping at the frame rate (2s) so that the impact of any interference is restricted to a single hop rather than degrading the entire transmission. Frequencies with persistent interference can be removed from the cell hopping sequence. Other mechanisms to remove interference include the base station directing antenna nulls towards strong sources of interference, careful scheduling of transmissions to terminals to avoid the frequencies where they perceive the strongest interference and the use of spreading to make the signal more resistant to interference when all these other techniques are insufficient.

Finally, M2M traffic is often characterised by very short messages, for example a 30-byte smart meter reading. The MAC protocol is designed to add minimal signalling overhead to such messages to avoid highly inefficient transmission. This is done through flexible small packets with highly optimised header information.

A global standards body – the Weightless SIG – has been established to take the Weightless standard and complete it as a royalty-free fully open standard.

6 Conclusions

The value in machines having wireless communications has long been understood and a large market predicted for many years. That this has not transpired has been because of the difficulty of meeting all the requirements within the constraints of the available radio spectrum. These constraints changed significantly with the advent of white space which provides near-perfect spectrum with free access. However, the combination of the unique and unusual nature of that access and the very different characteristics of machine traffic compared to human traffic means that using any existing standard is far from optimal. Hence, the need for a standard designed specifically for machine communications within whitespace. The Weightless standard is now under development and is scheduled for completion in early 2013.

7 References

[Gurney 2008] David Gurney, Greg Buchwald, Larry Ecklund, Steve Kuffner, John Grosspietsch (2008). "Geo-location database techniques for incumbent protection in the TV white space", , *DYSPAN 2008 -*

3rd IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks, Vol. 3, no. 1, October 2008, pp. 232-240

[Nekovee 2010] Maziar Nekovee (2010). "Cognitive radio access to TV white spaces: Spectrum opportunities, commercial applications and remaining technology challenges", *DYSPAN 2010 - 4th IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Vol. 4, no. 4, April 2010

[Ofcom] See the Ofcom consultation published at http://stakeholders.ofcom.org.uk/consultations/geolocation/

[FCC] See the US Notice of Proposed Rulemaking published at http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-301650A1.doc

[DVB] The "Blue Book" available from http://www.dvb.org/technology/standards/

[Wang 2011] Jianfeng Wang, Monisha Ghosh, Kiran Challapali (2011). "Emerging Cognitive Radio Applications: A Survey", *IEEE Communications Magazine*, Vol. 49, no. 3, March 2011

[Fitch 2011] Michael Fitch, Maziar Nekovee, Santosh Kawade, Keith Briggs, Richard MacKenzie (2011). "Wireless Service Provision in TV White Space with Cognitive Radio Technology: A Telecom Operator's Perspective and Experience", *IEEE Communications Magazine*, Vol. 49, no. 3, March 2011

Biography

William is CTO and one of the founding directors of Neul, a company developing the Weightless machine-to-machine technology and network, which was formed at the start of 2011.

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 12 books, 90 papers, and 18 patents. He is a Visiting Professor at Surrey University and DeMontfort University, a member of Ofcom's Spectrum Advisory Board (OSAB) and a Fellow of the Royal Academy of Engineering, the IEEE and the IET where he is a Vice President. 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. He can be contacted at william.webb@neul.com.