The Role of Networking Standards in Building the Internet of Things

William WEBB
CTO, Neul, Cambridge, United Kingdom

Abstract: The Internet of Things (IoT) promises widespread connection of sensors and devices such that smart cities, connected transport and other similar visions can be realised. However, to realise this a new wireless technology is required that is optimised for this particular application, providing very low cost, extremely long battery life and excellent coverage. This wireless technology needs to be an open global standard in order to foster a vibrant eco-system of key players. This paper discusses the issues involved in developing standards in the wireless telecommunications space before using a possible new IoT wireless standard as a case study to illustrate the issues with standardisation.

Key words: Internet of Things, wireless, standards, Weightless, machine-to-machine, M2M, ETSI, IEEE, white space, IPR.

■ The Internet of Things needs a wireless solution

The Internet of Things (IoT) is not a well defined concept but in essence is the idea that many machines, sensors and other electronic devices could be connected back to a core database that could process and act upon the information that they transmit. For example, in a smart city sensors could measure traffic levels and cause the alteration of traffic light timing patterns. Or they could measure the level of rubbish in dustbins resulting in optimised rubbish collection schedules. Many thousands of such applications could be envisaged.

There are four key components to such a system:

- Sensors that monitor some physical parameter such as temperature.
- A communications system that relays information from the sensor to a database.
- A central database and control system that analyses the information and sends control signals to some device.
- A control device, such as a traffic light.

Most of these are already widely available, or readily produced. Sensors are very low cost and miniaturised – for example mobile phones routinely contain a wide variety of them. Central databases and control systems have long been used for a variety of tasks and most control devices are already in place. The element that is lacking, and the reason why the IoT is not widely available, is a wireless solution that meets all the requirements for machine communications. While there are many possible applications the common characteristics tend to be (1) very low cost both for the chipset and the annual fee for sending data, (2) ubiquitous coverage, even better than cellular, (3) in some cases battery life of 10 years or so. These requirements are obviously challenging, but there are some characteristics that can be exploited in system design including (1) most messages are very short, (2) delays of a few seconds are rarely problematic, (3) data rates can be low, (4) sleep times can be long in some cases and (5) seamless handover is not needed.

Looking at this list it becomes obvious that there is no standard available that comes close to delivering this. Cellular provides coverage to almost the level needed but cannot achieve the cost points and battery life required. Indeed, if it could it would have done long ago. It is moving ever further away from a design that would exploit the characteristics that would provide such a system as it becomes ever faster and more complex. Bluetooth and Zigbee do not provide the range and coverage needed. Long-standing data systems like Paknet do not have the capacity or low device costs. Hence the need for a new wireless solution optimised for the IoT. However, introduction of a new solution is never simple because of the need for standards as the remainder of this paper discusses.

■ Wireless technologies have a "two ends" problem

Wireless systems differ from many other goods in that one wireless device is rarely of any use. It takes at least two to form a wireless link, and often thousands or millions to form a useful system. This is different from goods like cars, washing machines, computers, etc., which are equally valuable regardless of whether there are other similar devices around. In some simple cases an individual might buy both ends of the wireless link themselves at the same time – home cordless phones are a good example of this. But for most useful systems, one person will buy one device and a different person, or company, the other device. Even where one person buys all elements, such as a Bluetooth-enabled phone and a car with Bluetooth connectivity, they may buy them at different times from different vendors. For the IoT it is highly likely that devices will be sourced from a wide range of vendors and need to communicate with networks whose equipment is sourced from other vendors.

This "two ends" problem leads to a number of requirements:

- If systems are sourced from different vendors or at different times there must be a defined standard to ensure that they "interoperate".
- There is somewhat of a chicken-and-egg problem in that there is little incentive to be the person who buys the first end of the link. There may need to be some entity that acts to ease this.

For example, with Bluetooth there was firstly a defined standard to ensure interoperability. Next handset vendors embedded Bluetooth chipsets into billions of handsets despite there being no other Bluetooth enabled devices. This finally persuaded manufacturers of wireless headsets to use Bluetooth. As devices became more widespread many others were persuaded of the success of Bluetooth and built an increasing range of devices fuelling a virtuous circle.

The next section considers how standards can lead to such a virtuous circle and indeed why only standards deliver success in wireless.

■ The only wireless technologies that succeed are standards

If we examine wireless technologies we find that, almost without exception, there are no successful wireless technologies that are not open international standards. These include WiFi (IEEE 802.11), Bluetooth, 3G (and now 4G) and DECT cordless phones. Further, there is almost invariably only one successful global standard for each application area. There is only one standard for wireless networking (WiFi – albeit in many variants). There is only one for short-range personal connectivity – Bluetooth. There have been two standards for cellular, one from the US and one for the rest of the world but this is consolidating into a single standard with 4G. Where countries have tried to develop their own national wireless standards, such as Japan with 2G and 3G standards and China with 3G standards this has ended in failure as large vendors have opted not to develop equipment.

As well as overcoming the two-ends issue, successful standards bring other benefits, including:

- a competitive supply of equipment, reducing the risk for a purchasing company of being locked into a monopolistic supplier;
- economies of scale from large volume manufacturers, dramatically reducing device costs by making large-scale integration into complex chipsets worthwhile;
- the ability for devices to roam across countries and world-wide, for example enabling a WiFi enabled laptop to work in any country;
- greater innovation and inventiveness from bringing together an eco-system of interested players;
- often a better technology as a result of the peer-review that takes place as the standard is being developed (although see some caveats to this below).

History strongly indicates that only a standard for the IoT will succeed as opposed to a proprietary product. Further, there are many logical reasons for this. In the next few sections we consider how standards come about and some of the problems and issues with the standards process.

■ Traditional approaches to developing standards are failing

The classic view of standards is that a range of companies come together under the auspices of a standards body and work harmoniously to deliver the necessary standard. There have been cases in the past when this has happened – GSM is a reasonable example. However, competitive tensions tend to result in complexities, issues and sometimes divergent routes.

The tension arises from the desire of companies to profit from their work in the standard. The issue is akin to game theory. If all work harmoniously and equally divide the rewards then all benefit from rapid and successful development of the standard. But if a few key players manage to extract extraordinary returns then others may do better to develop a different standard. In the case of standards, this competitive game is played out through the vehicle of intellectual property rights (IPR). If a company holds IPR which becomes a core part of the standard then it can charge royalties to any other company that implements the standard. Such IPR is often termed "essential IPR" because use of it is essential in implementing the standard. Historically, those with essential IPR in key standards have achieved above average returns on investments. For example, Qualcomm hold many essential patents in 3G and 4G wireless technologies and have been more profitable than their peers as a result. Hence an incentive on those participating in standards to inject techniques into the standards process in which they have IPR. IPR is becoming ever more important in telecommunications as large companies such as Google and Microsoft buy up "patent pools" in order to battle against other established players around the world. At the time of writing devices were being removed from sale almost weekly in various countries due to alleged patent infringements. The optimal way to resolve these issues would be a change to patent rules, perhaps considering approaches such as compulsory patents, but despite major reviews in the US and elsewhere, significant change appears unlikely in the foreseeable future.

In some cases these tensions have become so severe they have derailed the standards process. For example, a range of companies came together to develop a standard for very short range, very high rate communications called ultra-wideband (UWB) back in the 2000 to around 2004. They worked together within the IEEE, a recognised standards body. However, there were two possible technical approaches to realising UWB and coalitions of companies with key IPR in each approach formed. They were unable to resolve these differences and as a result the standards process stalled. Eventually one coalition decided to form their own standards body to pursue their preferred solution. However, by this time many of those working in this space ran out of funding while those who might have adopted the solution held back from investment until there was certainty. The end result was that the initiative failed completely. Even where tensions can be accommodated they can lead to sub-optimal standards which either avoid useful technology or add in multiple unnecessary options in an effort to keep all key players on board.

There can also be competition between standards bodies. In the wireless arena there are two pre-eminent entities. These are the IEEE, a US body, and ETSI, a European body. In addition, there are standards bodies in China, Japan and other national entities. The reason that the US partly adopted a different standard for 2G and 3G was because of a preference towards their national standards body. To some degree this has been resolved in later variants of 3G and 4G through a consortium of standards bodies known as 3GPP, but this only applies to cellular systems and there is still the possibility of competition in areas such as the IoT where both ETSI and the IEEE are showing interest in starting a process of standardising a wireless solution.

Even where the standard process proceeds relatively smoothly, there are problems with conventional standards bodies. Often the bureaucracy involved and the large number of interested parties can result in the standards process taking three or more years. This has led some to seek a different approach. An example of this is Bluetooth.

In the case of Bluetooth the basic concept for the technology, for short peer-to-peer connectivity, already existed. Ericsson had developed a technology they called MCLink but had been unsuccessful in implementing it, precisely for the reasons cited above such as the two-ends problem and the need for a standard. Rather than approach a standards body they decided, along with Intel, Toshiba and others, to establish their own, with the sole aim of turning MCLink into an open standard. They called their new entity the Bluetooth Special Interest Group or SIG. A key decision was to make licensing of the key IPR royalty-free. This made it much more attractive for others to join who might otherwise have been suspicious that Ericsson would hold most key IPR and exert excessive power over the standard. Of course this was not so attractive for Ericsson, but for them Bluetooth was seen more as a vehicle to sell more mobile phones rather than to profit from directly. As is well known, the approach was very successful, leading to Bluetooth becoming one of the world's key wireless standards. The royalty-free regime of Bluetooth has led to far fewer IPR disputes than with other wireless technologies suggesting alternative IPR models may be a way around the current issues with the patent regime.

A further issue to consider is the increasing trend towards new wireless technologies being pioneered by small start-up companies. Over time, the larger entities such as Ericsson or Vodafone have cut back on their own internal R&D leaving it to start-ups to pioneer new ideas. This can be seen, for example, in the area of small "femtocells" where initially all the chip vendors and femtocell manufacturers were start-ups, with these slowly being acquired by larger players as the success of the concept became clearer. Start-ups do not work well within the conventional standards approach as the timescales are often longer than they can obtain funding for, the need for legal resources draining and the whole culture of the conventional standards approach somewhat counter to that of individuals within start-ups. Perhaps that is why there have been no completely new and successful wireless standards for around a decade now. Instead, we have seen further iterations of existing standards such as cellular and WiFi. Standards then are critical to the development of new wireless systems but the process seems to be failing, particularly for innovative new concepts.

In the next section we consider the background to the development of a technology for the IoT and then subsequently consider which approach should be adopted to develop a standard for the IoT.

■ Standards for the Internet of Things

Towards the beginning of this paper we noted that no standard currently existed that met the requirement for machine communications including extremely low cost, long battery life and ubiquitous coverage. We cannot achieve such a standard through modification of existing standards as these are optimised for personal communications which has completely different needs from machine communications. However, a standard is essential to enabling widespread deployment of sensors.

No large company has started to develop anything appropriate, instead there are a few start-up companies working in this space. One example of a new technology being developed is called "Weightless" [Webb 2012]. This takes advantage of an important development in radio spectrum management – the opening of a shared access band of spectrum termed "white space" (GURNEY, 2008; NEKOVEE, 2010) to deliver a technology specifically focussed on M2M connectivity.

White space spectrum is the unused portions of the spectrum band in and around TV transmissions. White space meets all of the requirements for IoT 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, 2010; FCC, 2010; DVB) but it is clear it will require low output power and stringent adjacent channel emission limits, necessitating bespoke radio design. Interference can be problematic in white space. Many channels have residual signals from TV transmissions. These can be emissions from distant, powerful TV masts that are too weak for useful TV reception but still significantly above the noise floor. 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. Hence, the need for a new standard.

■ The approach to standardisation used for the IoT

Drawing together the conclusions so far from this paper we can see that:

- A new wireless technology is required to enable the IoT.
- Wireless technologies only succeed if they are open global standards.
- The standards process has its issues, with the most successful recent example being Bluetooth.
- The IoT technology needs to be novel because of the use of white space and the particular requirements of machines.
- The structure of the industry suggests a technology is most likely to be pioneered by a start-up rather than a large established player.
- The conventional standards process does not fit well with start-ups.

Following this logic, Neul and other key players decided to take the Weightless technology into a global standard using the "Bluetooth approach". In 2011 the Weightless SIG was formed (Weightless), analogous to the Bluetooth SIG, as a vehicle not only to develop the Weightless standard for M2M communications but also to handle other important aspects such as testing and certification, marketing and promotion of the technology. As with Bluetooth, licensing for terminals is royalty free. It is hoped that this will overcome many of the patent problems currently besetting telecommunications standards, as discussed earlier. However, Weightless differs from Bluetooth in that Weightless requires a network communicating to devices, rather than devices talking directly to other devices. On the network side, Weightless has adopted the more conventional royalty approach used by most standards bodies known as "fair, reasonable and non-discriminatory (FRAND)" royalty payments. This allows some return to companies investing in R&D in this space but with few manufacturers of base stations expected, helps contain any risk. The Weightless SIG expects to complete its first version of the standard early in 2013 with devices becoming available towards the end of 2013.

At the time of writing in July 2012 the Weightless SIG was growing steadily with over 40 members. Opinion remained divided on whether M2M could be addressed using conventional cellular technology such as GPRS or would be better served by waiting for the completion of a bespoke standard. In some cases, hybrid solutions were proposed with GPRS carrying the traffic until the Weightless networks were sufficiently widespread. It seems likely that Weightless networks will be deployed during 2013 and 2014 at which point M2M application providers will be able to make an informed choice as to whether this new wireless standard is their best option for serving their communications needs.

■ Conclusions

In this paper we have shown that a new wireless technology designed specifically for machine applications is needed to enable the IoT and that such a technology does not currently exist. We have also discussed why this technology will need to be an open global standard. We then discussed the standards process, noting that it had an array of problems predominantly centred around tensions caused by IPR which resulted in extended timescales and even failure in some cases. While the existing standards bodies work well for extensions of major standards they do not appear well-suited to developing completely new standards, especially where the technology is being pioneered by start-ups or small companies. Instead, a different approach is being adopted of developing the standard in a bespoke entity established just for that purpose and easing IPR issues by adopting a royalty-free model especially for the large volume terminal devices. It may be that the wireless community, including standards bodies, regulators and industry, should look more generally at the standards development process and consider whether different approaches are needed in some cases.

References

DVB: *The Blue Book*. http://www.dvb.org/technology/standards/

FCC (2010): See the US Notice of Proposed Rulemaking published at: http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-301650A1.doc

GURNEY D., BUCHWALD G., ECKLUND L., KUFFNER S., GROSSPIETSCH J (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, pp. 232-240.

NEKOVEE M. (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.

Ofcom (2010): See the Ofcom consultation published at: http://stakeholders.ofcom.org.uk/consultations/geolocation/

WEBB W. (2012): *Understanding Weightless*, Cambridge University Press.

Weightless: see www.weightless.org