MOTT MACDONALD

The Internet of Things

An introduction

December 2016

10 Fleet Place
London EC4M 7RB
United Kingdom

T +44 (0)20 7651 0300
F +44 (0)20 7248 2698
mottmac.com

The Internet of Things

An introduction

December 2016

Information class: Standard

This document is issued for the party which commissioned it and for specific purposes connected with the above-captioned project only. It should not be relied upon by any other party or used for any other purpose.

We accept no responsibility for the consequences of this document being relied upon by any other party, or being used for any other purpose, or containing any error or omission which is due to an error or omission in data supplied to us by other parties.

This document contains confidential information and proprietary intellectual property. It should not be shown to other parties without consent from us and from the party which commissioned it.

Contents

Executive summary 1

1 Introduction 2

1.1 Architecture 2
1.2 Machine-to-Machine (M2M) Communications 3
1.3 Peer-to-Peer 4

2 Applications 5

2.1 Industry 5
2.2 Transport 5
2.3 Utilities 5
2.4 Smart Buildings & Infrastructure 6
2.5 Healthcare 7
2.6 Safety & Security 7
2.7 Tracking of Tagged Items 8

3 Communication Technologies 9

3.1 Requirements 9
3.2 Communication Technologies for End Devices 10
3.2.1 Cellular Networks 10
3.2.2 Licensed Spectrum 11
3.2.3 Unlicensed Spectrum 12
3.2.4 White Space 12

4 Application Platforms 15

5 Conclusions 17

5.1 Future Potential 17
5.2 Possible Problems 17

Executive summary

A recent magazine article described an A-to-Z of sensor applications:

Aerospace, Beverage vending, Crops, Disease detection, Earthquakes, Forest fires, Greenhouses, Health, Internet search, Jellyfish, Kitchens, Lighting, Medical devices, NFC payment, Odour, Parking, Quantum-cascade laser, Radiation monitoring, Sport, Traffic, Underground, Vehicle accident, Waste management, Xbox 360, Yachts, Zebras.

Sensors such as these will form the basis of the Internet of Things (IoT). Intelligent IoT applications are likely to affect almost every aspect of our lives.

IoT applications monitor and interpret their environment based on large amounts of incoming data from sensors and other data sources. This data is analysed in sophisticated ways that enable IoT applications to respond appropriately to external events that they may not have seen before. This means that IoT applications are not restricted to monitoring and interpretation - they can also be used to control things.

As can be seen from the A-to-Z list above, the same basic architecture is applicable to a very wide range of potential applications. Some of these are discussed in Section 2. According to Forrester Research, global IoT revenues will be thirty times those of the internet by 2020, making it the next trillion dollar communication industry.

This white paper provides an introduction to the Internet of Things. It describes how the IoT architecture is likely to evolve, and considers the future potential of the technology along with some possible problems.

1 Introduction

1.1 Architecture

At its heart, the Internet of Things consists of billions of monitoring and control devices that use the internet to send data back to a central point for analysis, as illustrated in Figure 1:

Figure 1: IoT Architecture

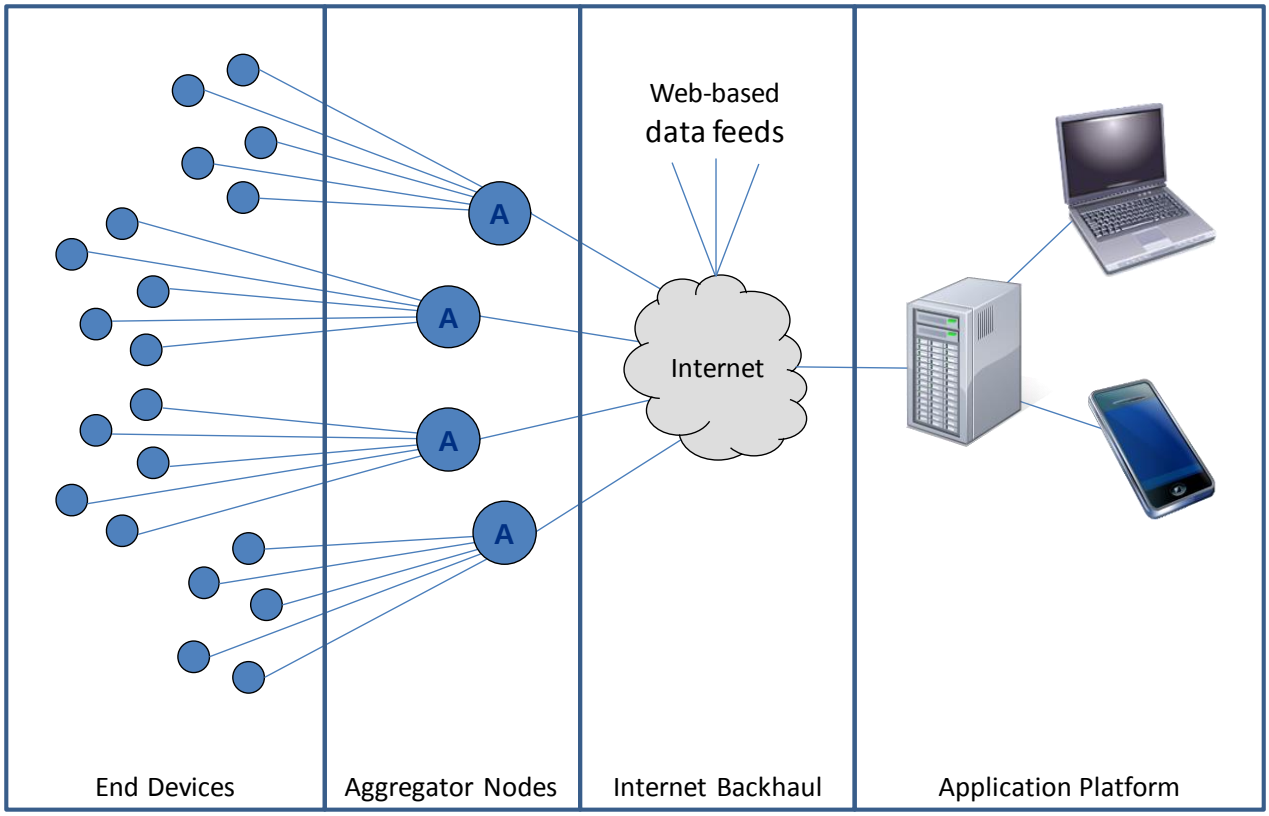

Source: Mott MacDonald

The End Devices shown in Figure 1 could include:

- Sensors that measure a physical quantity (eg temperature, pressure, vibration) or report on status (eg valve open / shut, alarm on / off)
- Actuators that enable some form of remote control (eg turning equipment on / off, changing the setting of a valve)

Although current end devices typically contain a significant level of local intelligence and communication capability, it is anticipated that the Internet of Things will evolve towards applications that use much larger numbers of simple, low-cost end devices. For example, a single home could contain as many as 1,500 devices to monitor and control everything from lighting and heating to entertainment and security systems. To make such a concept viable,

these devices would have to be very cheap to produce, consume very little power and communicate without the need for cables.

The primary role of the Aggregator Nodes shown in Figure 1 is to interface the end devices to the internet. Low-cost end devices cannot justify the cost or the power required for an internet connection, and connecting them directly to the internet would open up a range of security issues. Furthermore, these end devices typically generate very little data (maybe only a few bits per hour) and internet protocols would be hopelessly inefficient at such low bit rates. The aggregator nodes therefore enable a large number of very simple, low cost devices to communicate efficiently and securely across the internet. Where a small number of more sophisticated end devices are used, they might be connected directly to the internet without the need for an aggregator node.

The Application Platform shown in Figure 1 collects and manages the “big data” that can be generated by a very large array of end devices. It also provides sophisticated data analytics to extract useful information from the raw data, and to present this information on convenient user interfaces such as laptop computers and smartphones. Where appropriate, an IoT application can be empowered to make decisions and take actions without human intervention.

In some cases, the required data can be obtained from web-based data sources, thereby reducing the number of sensors that need to be deployed and managed. Examples of such sources might include news feeds, transport information and real-time rainfall data. It is also possible to scavenge or infer data; for example, measurement of electrical harmonics and how they change over time can be used to characterise machinery without the need for a large number of separate sensors. As another example, scanning Twitter for certain key words can provide utilities with an early indication that their customers are experiencing problems in a particular area.

1.2 Machine-to-Machine (M2M) Communications

The Internet of Things evolved from the earlier concept of Machine-to-Machine (M2M) communications. Although the two terms are often used interchangeably, there are some significant differences as set out in Table 1 below:

Table 1: Comparison of M2M and IoT

Machine-to-Machine (M2M)	Internet of Things (IoT)
M2M applications are typically based on intelligent modules embedded within a machine.	In addition to standard M2M end devices, IoT also accommodates inexpensive, low-power devices.
M2M solutions typically rely on point-to-point communication using either cellular or wired networks.	IoT solutions are likely to rely on very low cost wireless access networks with internet backhaul.
M2M data is typically processed in a dedicated software application that is targeted at point solutions in service management applications.	The end device data may be processed using big data analytics and other enterprise applications to improve overall business performance.
M2M systems typically generate relatively low-level data that is useful to technical staff.	IoT outputs can indicate higher-level trends, and so can also be useful to marketing and management staff.

Although the distinction between M2M and IoT inevitably becomes blurred in some situations, it is generally true to say that IoT collects data from a larger range of devices, and processes the data to extract useful information in a much more sophisticated and flexible way. As a result, the range of applications for IoT goes well beyond traditional M2M applications.

1.3 Peer-to-Peer

In the IoT architecture described above, the end devices can only communicate with an application platform. However, the term Internet of Things is sometimes used to refer to arrangements in which all the end devices are directly connected to the internet. In this situation, it is worth considering whether it would be more efficient for the end devices to communicate directly (peer-to-peer) rather than via the application platform.

Adopting a peer-to-peer architecture is superficially attractive because it would allow the Internet of Things to escape from the limitations of centralised command and control. However, there are a number of reasons why this would not be such a good idea:

- If the end devices have enough processing power to support peer-to-peer communication, they would be significantly more expensive and power-hungry than the devices described in Section 1.1 above. Whilst there can be requirements for devices of this type (eg video cameras) within the IoT, a more normal requirement is for much simpler devices. Minimising the cost and power requirements of the end devices enables a much larger array of parameters to be monitored, and this significantly improves the quality of the information produced.
- There is no obvious reason why most end devices would need to talk to each other. Why would a pressure sensor need to talk to a thermostat or a light switch?
- Keeping objects as simple as possible minimises the risk of hardware / software failures, configuration errors and technology obsolescence.
- Peer-to peer networking would require the devices to be directly visible on the internet, and this would make them more vulnerable to hacking.
- Maintaining a tree structure rather than peer-to-peer mesh makes network routing decisions much simpler, thereby promoting scalability.

Having said that, it is possible to envisage some level of peer-to-peer networking between aggregator nodes, and this might be needed in some real-time control applications where latency must be kept to a minimum. Peer-to-peer networking might also be appropriate between expensive, highly-capable end devices such as smart TVs and audio systems; Qualcom have developed the AllJoyn open-source, peer-to-peer networking platform for applications of this type.

2 Applications

As explained in Section 1, the Internet of Things can trace its origins back to the monitoring of critical infrastructure and machinery used in industrial processes. However, IoT has extended the original M2M concept to address a much wider range of applications, as illustrated by some of the examples given in this section.

2.1 Industry

Machines are often monitored to detect problems such as overheating or excessive vibration. IoT takes things a stage further by allowing a range of parameters to be tracked over time and analysed so that conditions that previously led to a breakdown can be recognised and detected when they occur again. This approach often allows intervention to take place before a fault actually occurs, thereby avoiding the cost and disruption caused by unplanned downtime. Further efficiency improvements can come from benchmarking the performance of a machine with comparable machines in other factories and by migrating to condition-based maintenance regimes.

The IoT has an important role to play in situations where plant and machinery is located in remote or inaccessible places. For example, it can be used for integrity monitoring of oil pipelines, storage tanks and pumps, and any damage or deterioration can be correlated with environmental factors such as weather and pollution.

2.2 Transport

Railways use track-mounted condition monitoring devices to check for train problems such as worn wheel rims and overheating bearings. Aircraft engine manufacturers offer services that monitor the performance of their engines in flight to detect potential problems before they become critical.

Major roads can be fitted with sensors to detect the volume of traffic and identify the location of any accidents. This information can be correlated with social media and mobile phone traffic in the area to enhance the quality of the information. Active signage can be used to adjust speed limits dynamically to suit weather and traffic conditions, and cars can connect to the system to obtain this information directly. IoT can also help motorists to find the nearest empty parking space in a town centre.

2.3 Utilities

IoT-based smart grid technologies are enabling power networks to reduce their carbon footprint and operate more efficiently. These technologies will become increasingly important as we move to a low-carbon world in which electricity displaces fossil fuels for transport (electric vehicles) and heating (heat pumps). Integrating IoT devices into anything that consumes electricity allows demand patterns to be monitored and modified to match the available supply. Demand management allows peaks in consumption to be flattened so that generation plant (including renewable generation) can be used more efficiently. IoT can also help to reduce repair times, improve network safety, reduce energy theft and provide much more accurate predictions of customers' future energy usage.

In water networks, IoT applications can improve the management of water pressure (and pressure transients) to reduce the incidence of burst pipes. They can also predict the impact that weather events will have on drains and sewage systems, thereby allowing action to be taken to minimise the risk of overflows. Energy costs for remotely-located plant can be optimised, and maintenance costs can be reduced by proactively detecting problems before they become critical. Many customer problems can be detected by monitoring Twitter for certain key words.

2.4 Smart Buildings & Infrastructure

Smart buildings attempt to adapt to their residents' changing needs and personal preferences during the course of a day. This can include adjusting the heating or air conditioning depending upon the weather forecast and whether anyone is at home. Lighting, curtains and music can all be controlled to make the home more welcoming or to deter burglars. Automatic settings can be manually overridden from a smartphone app, so a system that expects the occupant to return from work at a particular time (based on past experience) can be told if they will be staying out late.

The idea of buildings that adapt to their users can be extended to smart hotels, smart offices, smart shops and smart cities. These cities could include smart bridges, smart tunnels and smart public spaces – and every other conceivable form of smart infrastructure. An IoT application could monitor infrastructure for indications of environmental conditions or structural changes that could compromise safety, and might also schedule any necessary repairs or preventative maintenance.

The Chicago Array of Things provides an example of a public initiative to stimulate the development of IoT applications by deploying a wide range of sensors around a city and making the data freely available to application developers. Initiatives of this type are likely to stimulate economic activity in the areas where they occur, and so are likely to be widely copied[1].

[1] https://arrayofthings.github.io/

2.5 Healthcare

For some time now, the benefits of telehealth and telecare to both patients and the healthcare system have been widely recognised. Telehealth allows parameters such as heart rate, blood pressure or the performance of a pacemaker to be monitored remotely, thereby allowing patients with chronic conditions to be discharged from hospital. Telecare allows discreet monitoring of elderly or vulnerable people to enable them to live normal lives secure in the knowledge that help will arrive quickly if it is needed.

These trends have been reinforced by smartphone apps and “wearable tech” that allow medical issues to be identified early – often before the patient has become aware that anything is wrong.

2.6 Safety & Security

IoT end devices can include things like smoke alarms and intrusion detectors, thereby allowing the IoT to be used in safety and security applications. Since the IoT includes actuators as well as sensors, it is possible for the system to respond automatically to potentially-dangerous situations. For example, the appropriate response if smoke is detected might be to turn on the lights in the area, trigger a fire alarm, release the fire doors and summon the fire brigade.

Since an emergency can sometimes trigger large numbers of alarms that can confuse a human operator, the IoT should attempt to convert this raw data into useful information. This could be useful in earthquake or tsunami early-warning systems, where a large number of alarms can arise almost simultaneously.

2.7 Tracking of Tagged Items

Tracking the movement of tagged items is a key IoT capability. RFID tags are already used for tagging warehouse stock, parcels, cars, bicycles, pets, farm animals and even visitors to theme parks. IoT applications can also be used for fleet management and for tracking high-value assets.

3 Communication Technologies

The Internet of Things needs a detailed understanding of the environment that it is monitoring (and possibly controlling), so it is critically dependent on being able to communicate with a large number of sensors and other remote devices. This section considers the range of data communication technologies that are suitable for connecting the end devices to the aggregator nodes shown in Figure 1.

3.1 Requirements

Some key requirements for communicating with IoT end devices are set out in the table below:

Table 2: End Device Communication Requirements

	Requirements
Scalability	The number of end devices required by IoT applications is growing rapidly. In the near future, IoT communications solutions may have to handle billions of separate devices. It has been estimated that human beings in urban environments are each surrounded by 1,000 to 5,000 trackable objects[2], suggesting that the number of end devices globally will eventually be numbered in trillions.
Cost	In order to make the IoT concept viable for large numbers of end devices, the cost of each end device and its associated communication channel may need to be as low as a few pounds. For some potential applications, even this may prove to be too expensive, so end devices and the associated communication links need to be really simple and cheap. This almost certainly means that some form of wireless communication will be required to minimise installation and cabling costs. It also means that end devices must work straight out of the box without the need for setup and configuration procedures.
Working Life	There are some applications where the up-front cost of the sensor may not be the primary concern. Examples might include sensors that are embedded in infrastructure in a way that makes them hard or impossible to replace, such as strain gauges mounted inside suspension bridges or sensors in the core of a nuclear power station. In these circumstances, the working life of a sensor may have to match the working life of the infrastructure that it is monitoring, and the focus needs to be on whole-life cost rather than first-in cost.
Power	End devices connected by wireless links will often need to be battery-powered. Regular battery replacement is expensive, so devices such as smart gas meters are required to operate for at least ten years between battery changes. In some cases, end devices may be powered by energy harvesting rather than by a battery, but there would still be a need for very low power consumption.
Coverage	Many end devices are likely to be installed inside buildings – and some may even be in cellars or underground chambers – so radio links will require significant building penetration to avoid the need for site surveys. Some IoT applications will require national or even international coverage, but this will be provided by the internet backhaul.
Reliability	Communication with end devices must be reliable. Although the occasional loss of communication would not be a major problem, the inability to communicate with a large number of end devices as a result of a wider network problem is potentially more serious. Network problems could range from fading caused by heavy rain to a bug in a new software release.
Coexistence	For cost reasons, radio communication with end devices will normally have to take place in unlicensed frequency bands. The wireless technology selected will have to be able to share the spectrum with other users without causing or suffering from significant interference.
Security	As shown in Figure 1, IoT applications are normally connected to the internet so cyber security precautions are required to protect against threats such as hacking and malware. The IoT will also create new security issues, such as spoof sensors. Wherever possible, the

[2] Jean-Baptise Waldner, quoted in http://blog.nskinc.com/IT-Services-Boston/bid/206613/Apple-s-Continuity-and-the-Internet-of-Things

	Requirements
	aim should be to implement security overheads in the aggregator nodes rather than in the end devices.
Broadcast message capability	There are some situations in which the same information needs to be distributed to a large number of end devices at the same time. Sending a separate message to each device individually would be hopelessly inefficient, so a broadcast capability is required.
Small data packets	An IoT sensor might only need to transmit a few bytes of data per hour, but a standard IPv6 data packet requires a 40-byte header before it can transmit any data at all. The communication system must be able to handle very small data packets efficiently.
Real time capability	In some applications, a potentially-dangerous situation detected by an IoT sensor may require a rapid response. It must be possible for such a sensor to raise an alarm quickly without having to wait until it is polled.
Mobility	End devices can be located in cars, trains or aircraft, so mobility is sometimes required. Clearly, the need to support mobility will make the communications solution significantly more expensive, so the cost and power requirements listed above may have to be relaxed for mobile applications.
Standards	Open APIs will be needed to prevent the IoT from evolving as a series of “walled gardens”. Organisations involved in standardisation activities relevant to IoT include oneM2M, W3C, Open Geospatial Consortium (OGC), Industrial Internet Consortium, Home Gateway Initiative (HGI), IETF, ITU and IEEE.

3.2 Communication Technologies for End Devices

In some situations, end devices can be connected back to a central system using cables (either copper or fibre optic). Copper cables have the significant advantage that they can provide power to the end device as well as a data communication path. Copper cables are also likely to be simpler and cheaper to install than fibre optics, and the bandwidth available on copper cables is nearly always sufficient for IoT applications. However, fibre cables may be deployed in difficult or dangerous environments because they are not prone to electrical interference, do not present a fire risk and provide a higher level of data security.

Given the very large number of end devices that typically characterise an IoT application, radio is normally a more appropriate solution than cable for data communications. Possible radio technologies can be grouped into four main categories:

- Cellular networks
- Licensed spectrum
- Unlicensed spectrum
- White space

Some possible options are assessed in the following sections.

3.2.1 Cellular Networks

Cellular networks can be used to supply communication services for IoT applications. These would typically be GPRS services, but 3G or 4G mobile network services could be used for higher bandwidth requirements. In the UK, O2’s GPRS network will be used to connect smart meters in central and southern regions of the UK.

Some strengths and weaknesses of cellular networks for IoT applications are set out in the table below:

Table 3: Cellular Networks for IoT Applications

Advantages	Disadvantages
• Cellular networks already exist, so new applications can be implemented quickly. • Cellular networks offer reasonable coverage.	• Gaps in network coverage (particularly in-building coverage). For smart metering in some parts of the UK, O2's GPRS network will have to be supplemented by wireless mesh technology. • High power consumption. Smartphone users have become accustomed to recharging their devices almost every day, whilst IoT devices may be required to operate for ten years between battery changes. • High monthly charges. The commercial model for providing network connectivity to smartphone and tablet users is clearly very different from the commercial model for IoT devices. IoT applications would have to pay for bandwidth and mobility capabilities that they do not normally require. • Over time, the focus on providing smartphone users with higher bandwidth will move cellular network tariff models even further away from the low bandwidth / power / cost requirements of IoT. • The rapid pace of innovation in the mobile market means that older technologies such as GPRS could soon be withdrawn, leaving IoT users with an expensive upgrade path • Mobile network congestion can occur at peak times or during emergency situations. • Commercial mobile networks are not engineered to meet the demanding requirements of safety-critical applications.

3.2.2 Licensed Spectrum

One way to avoid interference from other radio users is to build a narrowband radio network using licensed spectrum. One example of this is the smart metering network that Arqiva will be building in Scotland and Northern England using Sensus Flexnet™ technology operating over licensed spectrum in the 400MHz band. Other examples are the Aclara® STAR® Network and TTP's Matrix.

Some strengths and weaknesses of using licensed spectrum for IoT applications are set out in the table below:

Table 4: Licensed Spectrum for IoT Applications

Advantages	Disadvantages
• The use of licensed spectrum means that the network should not experience interference from other radio users. • The use of a private network prevents traffic congestion caused by other network users. • Since the network is controlled by the user rather than by a commercial network operator, any necessary migration to a new networking technology can be timed to suit user requirements. • Licensed bands are likely to permit the use of higher power than unlicensed bands, leading to larger coverage areas and better building penetration.	• The need to obtain a spectrum license can delay (or even prevent) roll-out. • The spectrum license can add significant cost to an IoT application.

3.2.3 Unlicensed Spectrum

There are a number of suitable radio technologies for IoT applications that operate in unlicensed spectrum in the region of 868MHz in Europe (915MHz in the US) or 2.4GHz. As a result of the restricted transmit powers available in unlicensed bands, wireless mesh technologies are often used in which remote nodes communicate with the application platform by using adjacent nodes as repeaters. Examples of radio technologies that operate in unlicensed spectrum include:

- Silver Spring Networks Frequency Hopping Spread Spectrum Mesh (915 MHz)
- Itron Frequency Hopping Spread Spectrum Mesh (915 MHz)
- 802.15.4. ZigBee® Mesh (868MHz, 900MHz or 2.4GHz)
- 802.11s WiFi Mesh (2.4GHz, 5GHz)
- SigFox's Ultra Narrow Band (UNB) (868MHz in Europe; 915MHz in the US)
- On-Ramp's Total Reach Random Phase Multiple Access (RPMA) (2.4GHz)
- Telensa's PLANet street light control system (868MHz)
- TTP's Matrix (most major license-exempt bands)
- Semtech's LoRa (866MHz, 915MHz)
- WiFi-Direct (2.4GHz, 5GHz)

Some strengths and weaknesses of using unlicensed spectrum for IoT applications are set out in the table below:

Table 5: Unlicensed Spectrum for IoT Applications

Advantages	Disadvantages
• No need to obtain a spectrum license • No spectrum license fees • Many of these technologies are optimised for IoT or M2M applications. For example, Sigfox claim that communication with a smart meter using their network uses 100 times less power than if the communication went over a GSM network, thereby extending battery life from a few months to 20 years. • Wireless mesh topologies may provide a degree of resilience that is not available in simple point-to-multipoint topologies.	• When operating in unlicensed spectrum, there is always the risk of interference from other networks that are using the same spectrum. However, spread spectrum modulation techniques can be used to minimise the risk of interference. • Users in unlicensed bands have to operate at restricted power levels, so range is limited. This is not a problem in urban environments where the distance from one node to the next is relatively short, but can be a problem in rural areas. It can also be a problem during network roll-out, because distant nodes cannot communicate until intervening nodes have been built.

Arqiva is building a wireless network dedicated to supporting the Internet of Things across the UK. The network will initially go live in ten major cities across the UK, and will be based on SigFox's Ultra Narrow Band (UNB) radio technology.

3.2.4 White Space

"White Space" refers to the guard bands that are used to minimise interference between high-power UHF television transmitters. In spite of the fact that the white space has been deliberately left empty in order to prevent interference, it is possible to use this spectrum for other applications if certain rules are followed. For example:

- The equipment must transmit at low power, and use low gain antennas.
- Transmissions must use a modulation scheme that is unlikely to cause interference.
- The equipment needs access to GPS so that it can determine its location. Since television transmitters operate on different frequencies in different areas, and some frequencies are

only used at certain times of day, the equipment uses an online database to determine which frequencies are available locally at a particular time.

Subject to rules of this type, it is possible to open up white space sections of the TV broadcast spectrum to other users without causing interference problems. White space is already being used in countries such as UK, USA and Singapore.

Terrestrial television broadcasting lies in a valuable part of the spectrum because sub-800MHz frequencies can travel long distances and provide good penetration through foliage and into buildings. This means that white space is a particularly attractive option for Internet of Things applications. Of course, users still face the potential risk of interference from other users, but this is a problem in any unlicensed band, and sensing technology can be used to automatically select the least congested channel.

However, the lack of a single communications standard (such as WiFi in the ISM bands) creates new possibilities for interference. IEEE 802.11af (also referred to as White-Fi and Super Wi-Fi) is part of the 802.11 family of standards, and was approved in February 2014. It covers wireless local area network (WLAN) operation in TV white space spectrum in the VHF and UHF bands between 54 and 790 MHz.

Weightless is a white space technology that is specifically aimed at mainstream IoT applications. Operating distances range from a few metres to about 10 km, and data packets can be as small as 10 bytes.

SMS
SMS
SMS

4 Application Platforms

For reasons explained in Section 1.3, a peer-to-peer architecture is unlikely to be suitable for IoT applications. Instead, most of the intelligence in the system will be centralised in an Application Platform as shown in Figure 1. The primary functions of the application platform are:

- Data collection. The application platform will typically communicate with end devices indirectly via aggregator nodes. IP packets arriving at the application platform will therefore contain small amounts of data aggregated across a large number of end devices, so the application platform will have to extract this information and store it. It will also have to manage all the different locations from which it needs to obtain data: These could include private end devices that are only accessible by that specific application platform, public end devices that make their information available to multiple application platforms and web-based data feeds (such as weather forecasts) that can enable deeper meaning to be extracted from the data.
- Data segmentation. Once the data is stored in a database, then it should be possible to view the data in different ways. For example, if the application is monitoring energy usage across a global corporation, it may be necessary to group the data by sensor type, by geographical region or by business entity.
- Data analysis. One of the primary purposes of an IoT application is to extract useful information from the collected data. Although each end device is only likely to generate a small amount of data, the volume of end devices and other data sources that are likely to be used by IoT applications, and the long periods of time over which the data might be accumulated, suggests that “big data” analytics will be needed for some IoT applications. It will also be necessary to address issues with the quality of the data (eg spurious or missing readings).
- Control. In some cases, the role of the IoT application goes beyond assisting a human operator to make decisions and includes some level of delegated authority to make those decisions. This could be because the decision is relatively trivial, because the human operator is unlikely to be able to react fast enough, or because the IoT is actually likely to make a better decision. A good analogy would be the autopilot on a commercial aircraft, which can take instructions from the pilot but can also intervene if required.
- Human interfaces. The information extracted from the raw data needs to be presented via user-friendly interfaces that enable human operators to interact with it. These interfaces could be provided on laptop computers, tablets, smartphones or any other suitable device. It should also be possible to issue alerts via email, SMS or telephone call if the system detects that a critical situation may be developing.

An application platform is likely to run on standard off-the-shelf computing platforms, and standard techniques will be used to improve the resilience of the computing platform if this can be justified by the application. Whilst an IoT application to control energy consumption across a global corporation might require large amounts of servers and storage, controlling energy usage in a single home might require nothing more than a smartphone. In some cases, a distributed application platform might be used so that time-critical decisions can be made close to the point where they are needed, whilst higher-level decisions requiring a wider view can be made further

back in the network. Since an application platform is normally connected to the internet, its physical location is generally not important and it could be cloud-based.

There is not necessarily a one-to-one relationship between an end device (such as a sensor) and the application platform that handles the data that it generates. Although some end devices will be restricted to a specific application platform, others may publish their information to any application platform that wishes to subscribe to it. Application platforms should be able to collect data from a wide range of internet-based sources such as weather reports, TV schedules and details of public events. The information collected by an application platform might be related to a particular geographical area, but could instead be defined by a particular multi-national organisation or a particular market segment. Some IoT applications could have global reach.

A top-down approach to IoT application development will only take things so far – standards-based open platforms are required to allow real innovation to flourish. The plethora of highly original applications that have appeared on the worldwide web stand in stark contrast to the very limited range of network applications that were available in the days when network applications were controlled by the network operator. With this in mind, the Web of Things is an open architecture that provides an application layer for the Internet of Things in much the same way that the web provides an application layer for the internet. There are also a range of platforms (such as Thingworx, Raco Wireless's Omega DevCloud, Axeda's IoT Platform, Davra Networks' RuBAN and Dizmo's Interface of Things) that enable rich, interactive IoT applications to be built using "drag-and-drop" interfaces without the need for coding.

Some control applications work quite satisfactorily with relatively simple, pre-determined logic. For example, if an alarm clock is set for 6:30am, then it may be perfectly reasonable for a smart house to turn on the heating at 6am, start to raise the lighting level in the bedroom at 6:25, have fresh coffee ready for 7am and open the garage door at 7:15am. However, applications of this type assume that everyday life is predictable, and this is often not the case. An application that tries to help by doing the wrong thing is actually more of a hindrance than a help. Parents of small children will already be familiar with this type of "help".

Most IoT control applications are likely to go well beyond pre-defined logic and attempt to emulate human behaviour. They will try to learn how to respond appropriately in a range of possible scenarios so that they can deduce how to respond appropriately if they encounter a scenario that they have not seen before. In some cases, an IoT application might be able to identify useful data affinities that would be invisible to a human operator, so IoT applications may be permitted to select some of their own data sources rather than relying on human intervention to do this for them.

5 Conclusions

5.1 Future Potential

The Internet of Things has been described as “the Fourth Industrial Revolution”, and it seems likely that at least some of the hype will eventually be justified. However, as is often the case with disruptive new technologies, the initial ramp-up is relatively slow until a tipping point is reached and the market suddenly takes off. IoT has yet to produce a “killer application” that can drive the deployment of 100m+ IoT devices, but the groundwork required to reach this point (in terms of technical developments, supplier commitment and standardisation) is being put in place.

The new insights provided by IoT applications can make both natural and manmade systems more predictable, and this predictability can lead directly to more efficient operation and enhanced quality of service to end users. It is likely that the business case for IoT will often be based on these two key benefits. For example, the ability to move from reactive to proactive maintenance of machinery can significantly reduce operating costs and equipment downtime. Predictability also means that more automation can be introduced into processes, with a reduced requirement for operator intervention. Energy usage can be optimised, and resources can be used more efficiently.

One of the key differences between IoT applications and the M2M applications that preceded them is that IoT is expected to analyse much larger data sets in much more sophisticated ways, thereby providing information that is potentially far more useful. There are plenty of examples of organisations that already collect very large amounts of data but the information extracted is little more than a status report. There are also tantalising glimpses of what could be achieved by more sophisticated data processing. The collection of the data and the development of the data analytics need to advance hand-in-hand if the full potential of IoT is to be achieved.

5.2 Possible Problems

Although the Internet of Things appears to have a bright future, there are some potential problems that could prevent it from achieving its full potential. These problems have been divided into Technical and Societal problems in the tables below.

Table 6: Technical Problems

Problem	Details
Lack of Common Standards	As with many emerging technologies, IoT currently suffers from a lack of standards. If devices from one vendor cannot communicate with an application from another vendor, then potential applications will be severely restricted.
Lack of Network Addresses	Estimates of the number of devices connected to the Internet of Things vary considerably, but the number is already measured in billions and seems likely to reach 25-50 billion by the end of the decade. Not every device on the IoT will have its own IP address but, given the current shortage of IPv4 addresses on the internet, it seems inevitable that the global adoption of IPv6 (with its far larger range of addresses) will be critical for the successful development of IoT.
Managing Big Data	In recent years, the IT industry has placed considerable focus on the problems of managing and interpreting “big data”, so it is reasonable to assume that the technologies and processes required for this aspect of IoT have already been

Problem	Details
	developed. However, it is possible that some users of IoT applications will underestimate the quantity of data that they will have to manage.
Environmental Concerns	Adding electronics to mundane devices such as door handles or thermostats is a key feature of IoT. However, the electronics is likely to have a much shorter replacement cycle than the equipment in which it is embedded, leading to higher levels of waste. The recycling of semiconductors is a particular issue because of the toxic chemicals that are used during the manufacturing process.

Although potentially challenging, technical problems can typically be solved if consensus can be reached across a relatively small number of industry experts. Societal problems, on the other hand, tend to require wider public debate and have the potential to become political.

Table 7: Societal Problems

Problem	Details
Privacy	IoT applications can acquire a great deal of personal information that would be useful to 3rd parties for both legitimate and illegitimate reasons. For example, a smart home might know what time you are likely to leave the house on a particular day, and when you are likely to return. In fact, it might know a surprising amount about many of your personal habits, and may even be able to detect whether you are feeling unwell or unhappy. Some IoT applications will also use technologies such as RFID and GPS to track your location. Not surprisingly, privacy advocates and civil liberties campaigners have significant concerns about the potential use of IoT technologies by governments and large corporations to intrude into peoples' lives.
Security	As IoT applications migrate from interesting research projects to mission-critical business tools, the issue of security will become increasingly critical. IoT systems are likely to hold sensitive data relating to people and businesses, and many of the new applications proposed for IoT involve potentially-unsafe activities such as sharing data between systems and providing access to users via personal handheld devices. Furthermore, the IoT could be controlling dangerous machinery used in transport or industrial applications, so the risks are not confined to the virtual world. Although there is clear recognition that security is an issue, it seems inevitable that IoT applications will occasionally be hacked, and that real damage will occur as a result.
Legal Liability	If advanced IoT systems have the ability to initiate actions without the need for human intervention, then it raises some interesting legal questions about who is liable if they make a bad decision. Of course, industrial control systems have been taking potentially-dangerous decisions for many years, but the sophistication of some IoT systems is likely to go well beyond the deterministic logic used by most control systems. It is possible that current developments in driverless cars will help to clarify the law in this area.

Dr Andrew Wheen is a Project Director within Mott MacDonald's Digital Infrastructure consultancy. He is part of an integrated team of strategic advisers, technical experts and project managers that are delivering major telecommunications and IT projects around the world.

For further information, contact:

E-mail: Andrew.Wheen@mottmac.com
Office: +44 (0)20 7651 0656
Mott MacDonald Limited, 10 Fleet Place, London, EC4M 7RB, United Kingdom
www.mottmac.com