Going Beyond the User — The Challenges of Universal Connectivity in IoT

Darshana Thomas
Department of Electronic and Electrical Engineering
University of Strathclyde
Glasgow, G1 1XW
Email: darshana.thomas@strath.ac.uk

Greig Paul
Department of Electronic and Electrical Engineering
University of Strathclyde
Glasgow, G1 1XW
Email: greig.paul@strath.ac.uk

James Irvine
Department of Electronic and Electrical Engineering
University of Strathclyde
Glasgow, G1 1XW
Email: j.m.irvine@strath.ac.uk

***Abstract*—The Internet of Things (IoT) approach to interconnected devices has become a significant topic in recent years, and is likely to be a major influence on future networking standards, such as ongoing work on 5G. IoT introduces connectivity to a much wider range of devices than seen previously, which raises a number of challenges, both technical and ethical. This paper explores some of these challenges which IoT faces, as a result of the personal and confidential information which may be transmitted from body-worn sensors, and the inherent challenges of introducing connectivity to standalone devices, rather than to equipment operated by users.**

I. Introduction

The Internet of Things (IoT) is a networking paradigm encompassing ubiquitously connected devices, themselves no longer tied to a particular human user, which are capable of interaction with other such devices. IoT-based technologies are being investigated for health monitoring, smart homes and device control, vehicles and transportation, and in fixed-location physical deployments (such as within buildings or on street furniture) for environmental sensing.

Such diverse use-cases will inevitably encounter challenges, although we focus here on some specific challenges pertaining to the networking of such devices on the scales being considered, and factors to be considered when looking at carrying out such deployments. According to CISCO, 50 billion devices are expected to be connected to the internet by 2020 [1], raising questions over how connectivity for devices can be managed and provided. Such a significant rise in connected devices means significant changes in how we design and organise networks will be necessary.

IoT device communications can be considered to operate on one of three models. Firstly, a device might only transmit when a state-change is observed. For example, a remote temperature sensor may only transmit when the temperature changes, or falls outwith a permitted range. These devices will likely transmit sporadically, and would not necessarily require constant connectivity to carry out their activities. Of note, however, is the risk that in the event of an environmental (or other) change, many such devices may simultaneously attempt to establish connectivity, to report of a state change, potentially overwhelming the local access network. Similarly, in the case of devices on heavily contended access links, not receiving a status update is not in itself a guarantee of all being well — the device may be attempting to send a report, but may be unable to obtain reliable connectivity.

Secondly (to address the concern of this ambiguous *radio silence*), devices could be configured to report data on a scheduled basis. These devices will therefore only require connectivity at predictable periods, and the demand from these devices may be balanced. Means of automatically balancing this load (perhaps by allowing adjacent devices to observe patterns of usage, and select a period of relatively low activity for their own transmissions) offer a practical way to reduce the risk of connectivity problems. A hybrid of scheduled and conditional data transmission could be used, where abnormal conditions may be reported immediately, but regular data would only be reported on a scheduled basis. As such, a missed report may be detected and acted upon.

Finally, devices may also be designed to hold a constant connection to the internet, such that they may have their state queried in real-time, or receive (and react to) incoming requests. The requirement of a specific device will naturally define the connectivity mode, although requiring persistent inbound connectivity will require that the device remain associated with the access network at all times. If a network has a limit to the number of associated devices permitted (such as WiFi), this may pose a practical challenge during implementation.

Previous work has explored challenges of IoT [2], [3], although the majority of this work has focused on the the security aspects of IoT. In this extended abstract, we explore the challenges of achieving connectivity for IoT devices, which are inherently non-user oriented. We also explore some security and privacy considerations, relating specifically to the use of networked devices lacking human-in-the-loop input.

II. Deployment of IoT devices

When considering IoT deployments, the conventional model of services being consumed by users is no longer necessarily strictly accurate. Nonetheless, we shall consider that a connected device is, for our purposes, a user just like any other human user. Establishing quality-of-service guidelines to prioritise traffic of human users may well prove to be a stop-gap approach to prevent perceivable service quality degradation.

A. Cellular vs WiFi Connectivity

The number of users a mobile cell may sustain depends on the available bandwidth and resource blocks for the cell, and the quantities of data being transferred by users [4]. Cell types such as macrocells, femtocells and picocells play an important part in future IoT device connectivity. Based on the cell type and the number of users in the current cell, the ability to connect may be limited, particularly for larger cells. For a typical LTE macrocell, up to 300 users may be accommodated within a cell and have connectivity to the network [5]. Introducing new IoT devices, in addition to existing human users, may significantly raise the number of devices demanding connectivity within busy cells. This would result in increased traffic over the network.

Ongoing research into the specification and development of 5G will likely need to address the challenge of allocating limited resources to both mobile users and connected IoT devices. While upgrades to previous generation networks considered headline data speeds, we believe that priority should be given to increasing the number of users able to connect per cell, perhaps by reducing cell size, or implementing and deploying seamless hand-off between WiFi and cellular connections (to reduce the number of users dependent on one cell).

An alternative approach is to utilise existing WiFi networks for connected devices where possible. There are, however, also limitations on the number of WiFi users that can use the network per access point. Typically for a consumer-grade home wireless router, up to 10 users [6] can simultaneously use it. An enterprise access point featuring 2x2 MIMO can likely sustain up to 60 users [6].

When considering use of WiFi connectivity however, some extra considerations arise. While it is straightforward to connect a device to a single wireless network, the challenge of deploying access credentials for multiple networks, to multiple devices, is a very real one. To make WiFi conenctivity viable for connected devices, an out-of-band means of issuing new tokens to a disconnected device would be highly desireable, to avoid the need for manual intervention to install new keys on the device.

Similarly, however, the need to deploy SIM cards in cellular-connected devices may prove to be a challenge for implementations — each device requires its own SIM card, and changing cellular network would require a new SIM card to be installed on each device. Alternative means of handling the exchange of currently SIM-based subscriber information may be desirable, to remove the risk of expensive manual intervention being required on large numbers of deployed device.

Another approach is the aggregation of data from several IoT devices in a local concentrator (which has various means of connectivity available), which then relays these messages over its own links. While this does not resolve the challenge of wireless spectrum usage by the devices in a local setting, it would remove many of the challenges of deploying credentials to connected devices, if only the concentrator required updating. Wireless communication between devices and the concentrator would be over a much shorter distance however (akin to current domestic WiFi deployments), allowing for locally managed network access controls.

The cost of deploying cellular connectivity for these devices has to be considered by both vendors and end-users. Local data aggregation would help minimise the number of devices connected to a mobile network, although would increase the quantity of data transferred by each concentrator.

When considering cellular or WiFi connectivity of IoT devices, the relative mobility of devices should also be taken into account. Some devices operate in an inherently more fixed manner than others. For example, a climate sensor mounted on an item of street furniture (such as a lamp post) is much less likely to be mobile than an air pollution sensor located on a car-pool vehicle. Mobile devices will require more versatile and rapid hand-overs, compared to fixed devices (which may only need to handover in the event of a failure of a local access point). Devices requiring regular handovers will lead to greater resource usage on cellular networks.

B. Power consumption

Power consumption also needs to be considered when planning for deployments of connected devices. Aiming to utilise network capacity during off-peak periods would be one potential solution to reduce power consumption of IoT devices, as there would be less waiting required for a transmission slot, and less risk of collisions when attempting to associate with the network. It may not however always be practical to transmit during off-peak times, and if connected devices were to form the bulk of traffic, formerly off-peak periods may well become highly contended.

The distance of devices from the network base station is a significant consideration for connected device power consumption. For this reason, we believe that a focus on reducing cell size, and therefore maximum transmission distance, is an important consideration in the development of the 5G standard, and in planning for connected device deployments.

C. Resilience

It is possible that some deployments of connected devices will wish to ensure connectivity is sustained, even in the event of a failure of a local network. For example, in the event of a WiFi outage, or a power cut disrupting power to a local relaying device, the user may wish to ensure continued connectivity, at the expense of increased power usage, if the nature of the deployment demands that data continue to be available.

A key factor in resilience of connected devices (in additional to typical physical considerations, such as housing and protection against weather and abuse) will include the ability for the device to connect to new networks, or to select the most reliable network available. There are inherent considerations here for network providers, since users may wish to have the ability to migrate devices to a different network provider, if one proves unreliable, or a rival offers enhanced connectivity. Likewise, large deployments may wish to reduce their access costs by relaying data via concentrators (as discussed in Section II-A, or by prioritising local networks (such as WiFi) above cellular networks. This may open up new opportunities for one-providers connectivity solutions, offering both WiFi and cellular connectivity for devices using a single provider, or for marketplace-based approaches to connectivity [7], where

IoT device providers may negotiate bandwidth and access at scale in an open, competitive marketplace.

III. Security and Privacy

A device-oriented approach to networks, as opposed to the conventional user-oriented approach, presents a number of security and privacy considerations.

A. Security

Conventional security software and practices were often developed with the assumption that a user would be physically present at the system, in order to act as the decision-maker in the event of a question affecting data security. With an IoT-based device, where the device is no longer tied to an individual, with no human physically present, this model of security is no longer practical. For example, where a regular internet browser would prompt the user upon connection to a server with an invalid or expired SSL certificate, this is obviously not possible when considering user-less systems.

Previous work [8] has demonstrated the extreme risks of non-browser SSL certificate verification, and the extent to which widely deployed and relied-upon libraries and software failed to properly validate certificates. Without addressing these complexity issues, the security of data (and any remote updates) of IoT devices remains an open question.

As discussed in Section II-A, the secure distribution of access credentials for WiFi networks, and the management and issuance of SIM cards containing cellular network access credentials also pose a challenge for device operators. These form a security challenge as well, in that the secure bootstrapping of new devices may be a significant challenge — without out-the-box connectivity available to a device, provisioning of secure connectivity (and the secure exchange of access keys) will pose a significant challenge to device deployment.

B. Privacy

The concept of privacy in a device-oriented device is a new consideration (from the perspective of transmissions originating from the IoT device), although privacy is a long-established principle in the literature. Privacy, in this context referring to an individual's right and ability to control their personal or identifying information, ultimately refers to data belonging to an individual.

If an IoT device were to be designed (or maliciously modified, perhaps via an unauthorized firmware update) to monitor its network interfaces, and report on the identifiers of other devices, it would be possible for a network of IoT devices to become pervasive user-trackers. Previous works have considered these risks, within the context of traditional networks [9], although connected devices raise the risks of such tracking. Indeed, some of the proposals made in this work (from 2007) first saw wide deployment only recently, with Apple's launch of iOS 8 [10]. Even then, these anti-tracking features are only partially effective, with the device's original MAC address being restored when the device is in use. The technique of randomizing MAC address for scanning, when in the presence of pervasive networked devices, would likely become ineffective, given the ease with which addresses could be correlated as users make use of their devices.

C. Message Authentication

In order to prevent the introduction of invalid messages (for example, from a false sensor, or a device purporting to be a sensor), it is necessary for consumers of sensor data to verify the authenticity of readings and other communications received over any public or shared channel, such as wireless. The two key factors in verifying the authenticity of the message are the identity of the sending sensor node (via some form of device identifier or sensor number), and the integrity (and thus authenticity) of the data transmitted. This would be carried out using a digital signature or other data authentication mechanism. A more detailed discussion of specific protocols and techniques for sensor security is presented in [11].

Within conventional cellular networks, the concept of identity is managed by the Subscriber Identity Module (commonly known as a SIM card), as discussed in Section II-A. One key consideration here though, is that since SIMs are issued by the network operator, for use in connecting to their own network, the SIM-based identity of a wireless device (referred to as its IMSI) will change upon switching to a new network. Therefore, it is necessary to identify a sensor device using a hardware identifier (akin to the IMEI of cellular user equipment). In the past, IMEI identifiers have often suffered from *spoofing*, whereby they are cloned or altered. Therefore, we propose that strong authentication of device identity should be carried out, through the use of asymmetric keys. All data transmitted by a trusted sensor should be signed using a hardware-specific private key, allowing for its origin (the public key identity) to be verified. A malicious party wishing to present false data to the sensor network would therefore need to physically compromise the original sensor, exfiltrating the private keys, to generate correctly-signed future messages.

One further authentication challenge which is especially significant in wireless networks, is that of replay attacks. While a message may indeed present itself as originating from a valid sensor node on the network, and its contents may also be correctly signed to prove the authenticity of the data contained within, it may be an outdated message which is no longer relevant. When sensing critical data, a malicious party may intercept and store valid sensor readings from a time when sensor readings give no cause for concern, then re-transmit these during a time of elevated readings (in order to mask the problem). It is therefore essential to ensure that packets cannot successfully be replayed by an attacker with transmission capabilities, using a technique such as that discussed in [12]. More generically, each sensor reading should be authenticated, and contain a message counter field. Every message should increment the counter field, and the recipient should store the last received message counter value. In the event of a malicious party attempting to replay old readings, the decrease in counter would be visible, and an alert could be triggered to indicate the (unsuccessful) attempt to present fake readings. The ability to carry out a successful replay attack, however, depends upon the ability for an attacker to block the reception of the valid sensor readings, and transmit their own, which will be discussed in the next section.

IV. Conclusion

IoT devices face a number of challenges, in pursuit of the goal of universal connectivity for their devices, which this

extended abstract has explored. Key challenges include those of network capacity, and the cost of cellular network access (and the impact this may have on human subscribers), as well as power consumption, and the trade-offs between the use of mobile cellular networks, and fixed WiFi networks. A number of challenges with regard to security and privacy were also considered, specifically around the lack of a human user in-the-loop to make decisions on security issues, and the privacy implications of widespread sensor deployments with wireless network interfaces. We also introduced some proposals to resolve some of these challenges.

ACKNOWLEDGMENT

This work was partially funded by EPSRC Doctoral Training Grant EP/K503174/1.

REFERENCES

[1] R. Pepper. The Internet of Things is Now: M2M Devices Forecast 2013-2018. 2015.06.20. [Online]. Available: www.iicom.org/.../555-machine-to-machine-devices-forecast-2013-2018.

[2] V. Gazis, M. Goertz, M. Huber, A. Leonardi, K. Mathioudakis, A. Wiesmaier, and F. Zeiger, "Short paper: IoT: Challenges, projects, architectures," in *Intelligence in Next Generation Networks (ICIN), 2015 18th International Conference on*, Feb 2015, pp. 145–147.

[3] Z.-K. Zhang, M. Cho, C.-W. Wang, C.-W. Hsu, C.-K. Chen, and S. Shieh, "IoT security: Ongoing challenges and research opportunities," in *Service-Oriented Computing and Applications (SOCA), 2014 IEEE 7th International Conference on*, Nov 2014, pp. 230–234.

[4] D. Thomas and J. Irvine, "Connection and Resource Allocation of IoT sensors to cellular technology-LTE," in *Ph.D. Research in Microelectronics and Electronics (PRIME), 2015 11th Conference on*, June 2015, pp. 365–368.

[5] Trefor Davies. LTE the 4G deployment nitty gritty macrocells and small cells. 2015.07.20. [Online]. Available: http://www.trefor.net/2012/05/17/lte-the-4g-deployment-nitty-gritty-macrocells-and-small-cells/.

[6] How Many Users Can a Wireless Access Point Handle? 2015.07.20. [Online]. Available: http://www.securedgenetworks.com/blog/How-Many-Users-Can-a-Wireless-Access-Point-Handle.

[7] J. Irvine, "Adam Smith goes mobile: Managing services beyond 3G with the digital marketplace," *Invited Paper to European Wireless*, 2002.

[8] M. Georgiev, S. Iyengar, S. Jana, R. Anubhai, D. Boneh, and V. Shmatikov, "The most dangerous code in the world: Validating SSL certificates in non-browser software," in *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, ser. CCS '12. New York, NY, USA: ACM, 2012, pp. 38–49. [Online]. Available: http://doi.acm.org/10.1145/2382196.2382204

[9] T. Jiang, H. J. Wang, and Y.-C. Hu, "Preserving location privacy in wireless LANs," in *Proceedings of the 5th International Conference on Mobile Systems, Applications and Services*, ser. MobiSys '07. New York, NY, USA: ACM, 2007, pp. 246–257. [Online]. Available: http://doi.acm.org/10.1145/1247660.1247689

[10] L. Hutchinson. iOS 8 to stymie trackers and marketers with MAC address randomization. [Online]. Available: http://arstechnica.com/apple/2014/06/ios8-to-stymie-trackers-and-marketers-with-mac-address-randomization/

[11] P. Boyle and T. Newe, "Security protocols for use with wireless sensor networks: A survey of security architectures," in *Wireless and Mobile Communications, 2007. ICWMC '07. Third International Conference on*, March 2007, pp. 54–54.

[12] Z. Feng, J. Ning, I. Broustis, K. Pelechrinis, S. Krishnamurthy, and M. Faloutsos, "Coping with packet replay attacks in wireless networks," in *Sensor, Mesh and Ad Hoc Communications and Networks (SECON), 2011 8th Annual IEEE Communications Society Conference on*, June 2011, pp. 368–376.